

Oracle Database 19c to Amazon Aurora MySQL Migration Playbook

Oracle to Aurora MySQL Migration Playbook



Oracle to Aurora MySQL Migration Playbook: Oracle Database 19c to Amazon Aurora MySQL Migration Playbook

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Services or capabilities described in Amazon Web Services documentation might vary by Region. To see the differences applicable to the China Regions, see [Getting Started with Amazon Web Services in China \(PDF\)](#).

Table of Contents

Migration playbook from Oracle to Amazon Aurora MySQL overview	1
Feature compatibility tables	2
Feature compatibility legend	2
Amazon SCT and Amazon DMS automation level legend	3
Migration tools and services	5
Amazon Schema Conversion Tool	5
Download the software and drivers	6
Configure Amazon SCT	6
Create a new migration project	7
Amazon SCT action code index	10
Creating tables	11
Constraints	13
Data types	14
Common table expressions	15
Cursors	15
Transaction isolation	17
Stored procedures	18
Triggers	19
Sequences	21
Date and time functions	21
User-defined types	22
Synonyms	23
XML	23
MERGE	24
Query hints	24
Indexes	24
Partitioning	25
Materialized views	26
Views	26
UTL_Mail and UTL_SMTP	27
Database Links	28
PLSQL	28
EXECUTE IMMEDIATE	31
DBMS_OUTPUT	31

Amazon Database Migration Service	32
Migration tasks performed by Amazon DMS	32
How Amazon DMS works	33
Amazon RDS on Outposts	34
How it works	35
Amazon RDS Proxy	35
Amazon RDS Proxy benefits	36
How Amazon RDS Proxy works	37
Amazon Aurora Serverless v1	37
Amazon Aurora Serverless v2	39
How to provision	40
Amazon Aurora Parallel Query	42
Features	42
Benefits of Using Parallel Query	43
Important notes	43
Amazon Aurora Backtrack	43
Backtrack window	46
Backtracking limitations	47
SQL and PL/SQL	49
Single-row and aggregate Oracle and MySQL functions	50
Oracle usage	50
MySQL usage	51
Oracle and MySQL CREATE TABLE AS SELECT statement	63
Oracle usage	64
MySQL usage	64
Oracle and MySQL Common Table Expressions	65
Oracle usage	65
MySQL usage	66
Summary	70
Oracle sequences and identity columns and MySQL sequences and AUTO INCREMENT columns	70
Oracle usage	71
MySQL usage	75
Summary	78
Oracle and MySQL INSERT FROM SELECT statement	80
Oracle usage	81

MySQL usage	82
Multi-Version Concurrency Control	83
Oracle usage	84
MySQL usage	88
Summary	92
Oracle MERGE statement and MySQL equivalent	93
Oracle usage	93
MySQL usage	94
Summary	97
Oracle OLAP functions and MySQL Window functions	98
Oracle usage	99
MySQL usage	100
Oracle Transaction Model and MySQL Transactions	102
Oracle usage	102
MySQL usage	105
Summary	108
Oracle anonymous block and MySQL transactions or procedures	110
Oracle usage	111
MySQL usage	112
Conversion functions	113
Oracle usage	113
MySQL usage	118
Oracle and MySQL cursors	118
Oracle usage	118
MySQL usage	120
Summary	123
Oracle DBMS_DATAPUMP and MySQL integration with Amazon S3	125
Oracle usage	126
MySQL usage	128
Summary	128
Oracle DBMS_OUTPUT and MySQL SELECT	130
Oracle usage	130
MySQL usage	131
Oracle DBMS_RANDOM and MySQL RAND function	132
Oracle usage	132
MySQL usage	134

Oracle DBMS_REDEFINITION and MySQL tables and triggers	135
Oracle usage	135
MySQL usage	136
Oracle DBMS_SQL	136
Oracle usage	137
MySQL usage	139
Oracle EXECUTE IMMEDIATE and MySQL EXECUTE and PREPARE statements	139
Oracle usage	139
MySQL usage	141
Summary	142
Oracle procedures and functions and MySQL stored procedures	143
Oracle usage	143
MySQL usage	147
Summary	149
Oracle and MySQL regular expressions	150
Oracle usage	151
MySQL usage	152
Summary	154
Oracle TIMEZONE data type and functions and MySQL CONVERT_TZ function	154
Oracle usage	155
MySQL usage	157
Oracle and MySQL user-defined functions	158
Oracle usage	159
MySQL usage	160
Summary	162
Oracle UTL_FILE and MySQL integration with Amazon S3	163
Oracle usage	163
MySQL usage	164
Oracle UTL_MAIL or UTL_SMTP and Amazon Simple Notification Service	170
Oracle UTL_MAIL usage	170
Oracle UTL_SMTP usage	171
MySQL usage	172
Tables and indexes	175
Case sensitivity differences for Oracle and MySQL	175
Data types	177
Oracle usage	177

Oracle data types and MySQL data types	177
MySQL usage	184
Migration of Oracle data types to MySQL data types	188
Oracle Read-only tables and partitions and Amazon Aurora MySQL replicas	191
Oracle usage	191
MySQL usage	192
Oracle and MySQL table constraints	193
Oracle usage	194
MySQL usage	202
Oracle and MySQL temporary tables	209
Oracle usage	210
MySQL usage	212
Summary	213
Oracle and MySQL triggers	214
Oracle usage	215
MySQL usage	217
Summary	218
Oracle and MySQL tablespaces and data files	221
Oracle usage	221
MySQL usage	223
Summary	224
Oracle user-defined types	226
Oracle usage	227
Examples	227
MySQL usage	228
Oracle unused columns	228
Oracle usage	229
MySQL usage	230
Oracle virtual columns and MySQL generated columns	230
Oracle usage	230
MySQL usage	232
MySQL overall indexes summary	233
Usage	233
CREATE INDEX synopsis	234
Examples	235
Summary	236

Oracle bitmap indexes	237
Oracle usage	238
MySQL usage	238
Oracle and MySQL B-tree indexes	238
Oracle usage	239
MySQL usage	239
Example	240
Oracle composite indexes and MySQL multiple-column indexes	240
Oracle usage	241
MySQL usage	241
Oracle function-based indexes and MySQL indexing on generated columns	242
Oracle usage	242
MySQL usage	243
Oracle and MySQL invisible indexes	245
Oracle usage	246
MySQL usage	247
Oracle index-organized table and MySQL InnoDB clustered index	247
Oracle usage	248
MySQL usage	249
Oracle local and global partitioned indexes and MySQL partitioned indexes	250
Oracle usage	250
MySQL usage	251
Oracle automatic indexing	252
Oracle usage	252
MySQL usage	254
Special features and future content	255
Oracle Advanced Queuing and MySQL integration with Lambda	255
Oracle usage	256
MySQL usage	257
Oracle and MySQL character sets	257
Oracle usage	258
MySQL usage	259
Summary	260
Oracle database links and MySQL fully-qualified table names	261
Oracle usage	262
MySQL usage	263

Oracle DBMS_SCHEDULER and MySQL events	263
Oracle usage	264
MySQL usage	268
Summary	269
Oracle external tables and MySQL integration with Amazon S3	274
Oracle usage	275
MySQL usage	276
Oracle and MySQL inline views	282
Oracle usage	283
MySQL usage	283
Oracle JSON document support and MySQL JSON	284
Oracle usage	284
MySQL usage	285
Summary	289
Oracle materialized views and MySQL summary tables or views	290
Oracle usage	290
MySQL usage	293
Oracle multitenant and MySQL databases	293
Oracle usage	294
MySQL usage	297
Oracle Resource Manager and dedicated Amazon Aurora MySQL clusters	301
Oracle usage	301
MySQL usage	303
Summary	306
Oracle SecureFile LOBs and MySQL large objects	307
Oracle usage	308
MySQL usage	309
Oracle synonyms	310
Oracle usage	310
MySQL usage	312
Oracle and MySQL views	312
Oracle usage	312
MySQL usage	314
Oracle XML DB and MySQL XML	318
Oracle usage	318
MySQL usage	324

Summary	325
Oracle table compression	328
Oracle usage	329
MySQL usage	330
Oracle Log Miner and MySQL logs	330
Oracle usage	331
MySQL usage	332
Oracle SQL Result Cache and MySQL Query Cache	337
Oracle usage	337
MySQL usage	338
High availability and disaster recovery	340
Oracle Active Data Guard and MySQL replicas	340
Oracle usage	341
MySQL usage	342
Oracle Real Application Clusters and Aurora MySQL architecture	344
Oracle usage	344
MySQL usage	347
Summary	350
Migrate to Aurora MySQL Serverless	354
How it works	355
Oracle Traffic Director and Amazon RDS Proxy for Amazon Aurora MySQL	356
Oracle usage	357
MySQL usage	357
Oracle Data Pump and MySQL mysqldump and mysql	358
Oracle usage	358
MySQL usage	359
Summary	361
Oracle Flashback Database and MySQL snapshots	362
Oracle usage	363
MySQL usage	364
Summary	368
Oracle Flashback Table and MySQL snapshots	370
Oracle usage	371
MySQL usage	371
Summary	372
Oracle Recovery Manager and Amazon RDS snapshots	375

Oracle usage	375
MySQL usage	377
Summary	378
Oracle SQL*Loader and MySQL mysqlimport and LOAD DATA	383
Oracle usage	384
MySQL usage	385
Configuration	387
Oracle and Aurora MySQL upgrades	387
Oracle usage	387
MySQL usage	389
Upgrade using the Amazon Management Console	392
Upgrade using Amazon CLI	392
Summary	393
Oracle alert log and MySQL error log	394
Oracle usage	395
MySQL usage	396
Oracle SGA and PGA memory sizing and MySQL memory buffers	399
Oracle usage	400
MySQL usage	401
Summary	403
Oracle instance parameters and Aurora MySQL parameter groups	404
Oracle usage	405
MySQL usage	406
Oracle session parameters and MySQL session variables	407
Oracle usage	408
MySQL usage	409
Oracle and MySQL session parameter examples	410
Performance tuning	413
Database hints	413
Oracle usage	413
MySQL usage	414
Summary	418
Oracle and MySQL run plans	419
Oracle usage	419
MySQL usage	421
Oracle table statistics and MySQL managing statistics	424

Oracle usage	424
MySQL usage	426
Summary	428
Security	430
Encrypted connections	430
Oracle usage	430
MySQL usage	431
Oracle transparent data encryption and Amazon Aurora MySQL encryption and column encryption	431
Oracle usage	432
MySQL usage	436
Oracle roles and MySQL privileges	440
Oracle usage	440
MySQL usage	442
Oracle database users and MySQL users	443
Oracle usage	443
MySQL usage	445
Physical storage	449
Oracle and MySQL table partitioning	449
Oracle usage	449
MySQL Usage	454
Summary	458
Oracle sharding	459
Oracle usage	460
MySQL usage	460
Monitoring	461
Oracle usage	461
MySQL usage	462
Information schema tables	463
SHOW command	463
Summary	463
Migration quick tips	465
Management	465
SQL	465

Migration playbook from Oracle to Amazon Aurora MySQL overview

The first section of this document provides an overview of Amazon Schema Conversion Tool (Amazon SCT) and Amazon Database Migration Service (Amazon DMS) tools for automating the migration of schema, objects and data. The remainder of the document contains individual sections for the source database features and their Aurora counterparts. Each section provides a short overview of the feature, examples, and potential workaround solutions for incompatibilities.

You can use this playbook either as a reference to investigate the individual action codes generated by Amazon SCT, or to explore a variety of topics where you expect to have some incompatibility issues. When you use Amazon SCT, you may see a report that lists Action codes, which indicates some manual conversion is required, or that a manual verification is recommended. For your convenience, this Playbook includes an Amazon SCT Action Code Index section providing direct links to the relevant topics that discuss the manual conversion tasks needed to address these action codes. Alternatively, you can explore the Tables of Feature Compatibility section that provides high-level graphical indicators and descriptions of the feature compatibility between the source database and Aurora. It also includes a graphical compatibility indicator and links to the actual sections in the playbook.

The migration quick tips section provides a list of tips for administrators or developers who have little experience with Aurora (PostgreSQL or MySQL). It briefly highlights key differences between the source database and Aurora that they are likely to encounter.

Note that not all of the source database features are fully compatible with Aurora or have simple workarounds. From a migration perspective, this document doesn't yet cover all source database features and capabilities.

This database migration playbook covers the following topics:

- [Migration tools and services](#)
- [SQL and PL/SQL](#)
- [Special features and future content](#)
- [High availability and disaster recovery](#)
- [Configuration](#)
- [Performance tuning](#)

- [Security](#)
- [Storage](#)
- [Monitoring](#)
- [Migration quick tips](#)

Disclaimer

The various code snippets, commands, guides, best practices, and scripts included in this document should be used for reference only and are provided as-is without warranty. Test all of the code, commands, best practices, and scripts outlined in this document in a non-production environment first. Amazon and its affiliates are not responsible for any direct or indirect damage that may occur from the information contained in this document.

Feature compatibility tables

With Amazon DMS, you can ensure compatibility between the source and target databases during migration. Feature Compatibility defines the set of database engine features that Amazon DMS supports for a specific source-target combination. The following tables provide legends for feature compatibility to help you plan for your specific migration scenario.

Feature compatibility legend

Automation level icon	Description
	Very high compatibility. None or minimal low-risk and low-effort rewrites needed.
	High compatibility. Some low-risk rewrites needed, easy workarounds exist for incompatible features.
	Medium compatibility. More involved low-medium risk rewrites needed, some redesign may be needed for incompatible features.

Automation level icon	Description
	Low compatibility. Medium to high risk rewrites needed, some incompatible features require redesign and reasonable-effort workarounds exist.
	Very low compatibility. High risk and/or high-effort rewrites needed, some features require redesign and workarounds are challenging.
	Not compatible. No practical workarounds yet, may require an application level architectural solution to work around incompatibilities.

Amazon SCT and Amazon DMS automation level legend

Automation level icon	Description
	Full automation. Amazon SCT performs fully automatic conversion, no manual conversion needed.
	High automation. Minor, simple manual conversions may be needed.
	Medium automation. Low-medium complexity manual conversions may be needed.
	Low automation. Medium-high complexity manual conversions may be needed.

Automation level icon	Description
	Very low automation. High risk or complex manual conversions may be needed.
	No automation. Not currently supported by Amazon SCT, manual conversion is required for this feature.

Migration tools and services

Each of the pages in this section describe the various tools for automating the migration of schema, objects and data, and how to use them.

Topics

- [Amazon Schema Conversion Tool](#)
- [Amazon SCT action code index](#)
- [Amazon Database Migration Service](#)
- [Amazon RDS on Outposts](#)
- [Amazon RDS Proxy](#)
- [Amazon Aurora Serverless v1](#)
- [Amazon Aurora Parallel Query](#)
- [Amazon Aurora Backtrack](#)

Amazon Schema Conversion Tool

The Amazon Schema Conversion Tool (Amazon SCT) is a Java utility that connects to source and target databases, scans the source database schema objects (tables, views, indexes, procedures, and so on), and converts them to target database objects.

This section provides a step-by-step process for using Amazon SCT to migrate an Oracle database to an Aurora MySQL database cluster. Since Amazon SCT can automatically migrate most of the database objects, it greatly reduces manual effort.

We recommend to start every migration with the process outlined in this section and then use the rest of the Playbook to further explore manual solutions for objects that couldn't be migrated automatically. For more information, see the Amazon Schema Conversion Tool [User Guide](#).

Note

This walkthrough uses the Amazon Database Migration Service Sample Database. You can download it from [GitHub](#).

Download the software and drivers


Download and install Amazon SCT. For more information, see [Installing, verifying, and updating](#) in the Amazon Schema Conversion Tool User Guide.

Download the [Oracle](#) and [MySQL](#) drivers. For more information, see [Installing the required database drivers](#).

Configure Amazon SCT

Follow this procedure for configuring AWSSCT to streamline your database migration process.

1. Start Amazon Schema Conversion Tool (Amazon SCT).
2. Choose **Settings** and then choose **Global settings**.
3. On the left navigation bar, choose **Drivers**.
4. Enter the paths for the Oracle and MySQL drivers downloaded in the first step.



5. Choose **Apply** and then **OK**.

Create a new migration project

Create a new migration project to define the source and target databases, configure migration settings, and launch the replication process.

1. In Amazon SCT, choose **File**, and then choose **New project wizard**. Alternatively, use the keyboard shortcut **Ctrl+W**.
2. Enter a project name and select a location for the project files. For **Source engine**, choose **Oracle**, and then choose **Next**.
3. Enter connection details for the source Oracle database and choose **Test connection** to verify. Choose **Next**.

4. Select the schema or database to migrate and choose **Next**.
5. The progress bar displays the objects that Amazon SCT analyzes. When Amazon SCT completes the analysis, the application displays the database migration assessment report. Read the Executive summary and other sections. Note that the information on the screen is only partial. To read the full report, including details of the individual issues, choose **Save to PDF** at the top right and open the PDF document.

Create a new database migration project


Step 1. Choose a source
Step 2. Connect to the source database
Step 3. Choose a schema
Step 4. Run the database migration assessment
Step 5. Choose a target

Database Switch Assessment

Target platform	Auto or minimal changes			Complex actions			
	Storage objects	Code objects	Conversion actions	Storage objects		Code objects	
				Objects count	Conversion actions	Objects count	Conversion actions
Amazon RDS for MySQL	36 (100%)	11 (100%)	4	0 (0%)	0	0 (0%)	0
Amazon Aurora (MySQL compatible)	36 (100%)	11 (100%)	4	0 (0%)	0	0 (0%)	0
Amazon RDS for PostareSQL	36 (100%)	11 (100%)	2	0 (0%)	0	0 (0%)	0
Amazon Aurora (PostareSQL compatible)	36 (100%)	11 (100%)	2	0 (0%)	0	0 (0%)	0
Amazon RDS for MariaDB	36 (100%)	11 (100%)	4	0 (0%)	0	0 (0%)	0
Amazon Redshift	32 (89%)	6 (55%)	32	4 (11%)	2	5 (45%)	4
Amazon Glue	0 (0%)	9 (100%)	0	0 (0%)	0	0 (0%)	0

Previous Next Cancel

6. Scroll down to the **Database objects with conversion actions for Amazon Aurora (MySQL compatible)** section.



7. Scroll further down to the **Detailed recommendations for Amazon Aurora (MySQL compatible) migrations** section and review the migration recommendations.
8. Return to Amazon SCT and choose **Next**. Enter the connection details for the target Aurora MySQL database and choose **Finish**.
9. When the connection is complete, Amazon SCT displays the main window. In this interface, you can explore the individual issues and recommendations discovered by Amazon SCT.
10. Choose the schema, open the context (right-click) menu, and then choose **Create report** to create a report tailored for the target database type. You can view this report in Amazon SCT.
11. The progress bar updates while the report is generated.
12. Amazon SCT displays the executive summary page of the database migration assessment report.
13. Choose **Action items**. In this window, you can investigate each issue in detail and view the suggested course of action. For each issue, drill down to view all instances of that issue.
14. Choose the database name, open the context (right-click) menu, and choose **Convert schema**. Make sure that you uncheck the `sys` and `information_schema` system schemas. This step doesn't make any changes to the target database.

15. On the right pane, Amazon SCT displays the new virtual schema as if it exists in the target database. Drilling down into individual objects displays the actual syntax generated by Amazon SCT to migrate the objects.
16. Choose the database on the right pane, open the context (right-click) menu, and choose either **Apply to database** to automatically run the conversion script against the target database, or choose **Save as SQL** to save to an SQL file.
17. We recommend saving to an SQL file because you can verify and QA the converted code. Also, you can make the adjustments needed for objects that couldn't be automatically converted.

For more information, see the Amazon Schema Conversion Tool [User Guide](#).

Amazon SCT action code index

The following table shows the icons we use to describe the automation levels of Amazon Schema Conversion Tool (Amazon SCT) and Amazon Database Migration Service (Amazon DMS).

Automation level icon	Description
	Full automation — Amazon SCT performs fully automatic conversion, no manual conversion needed.
	High automation — Minor, simple manual conversions may be needed.
	Medium automation — Low-medium complexity manual conversions may be needed.
	Low automation — Medium-high complexity manual conversions may be needed.

Automation level icon	Description
	<p>Very low automation — High risk or complex manual conversions may be needed.</p>
	<p>No automation — Not currently supported by Amazon SCT, manual conversion is required for this feature.</p>

The following sections list the Amazon Schema Conversion Tool action codes for topics that are covered in this playbook.

Note

The links in the table point to the Oracle topic pages, which are immediately followed by the MySQL pages for the same topics.

Creating tables



Amazon SCT automatically converts the most commonly used constructs of the `CREATE TABLE` statement because Oracle and Amazon Aurora MySQL-Compatible Edition (Aurora MySQL) support the entry level American National Standards Institute (ANSI) compliance. These items include table names, containing security schema or database, column names, basic column data types, column and table constraints, column default values, primary, `UNIQUE`, and foreign keys. Some changes may be required for computed columns and global temporary tables.

Action code	Action message
73	MySQL doesn't support the <code>IDENTITY</code> statement with the <code>MAXVALUE</code> , <code>MINVALUE</code> , or

Action code	Action message
	CYCLE options or with the INCREMENT BY value that is different from 1.
74	MySQL doesn't support AUTO_INCREMENT statements without the primary key option on the same column.
190	MySQL doesn't support the COLUMN_VALUE pseudocolumn.
191	MySQL doesn't support the OBJECT_ID pseudocolumn.
192	MySQL doesn't support the ORA_ROWSCN pseudocolumn.
193	MySQL doesn't support the ROWID pseudocolumn.
198	MySQL doesn't support global temporary tables.
199	MySQL doesn't support clustered tables.
200	MySQL doesn't support external tables.
209	Amazon SCT uses triggers to emulate virtual columns because MySQL doesn't support virtual columns.
210	Amazon SCT uses triggers to emulate the usage of functions or expressions as default value in CREATE TABLE statements.
215	MySQL doesn't support virtual columns with unsupported build-in functions.

Action code	Action message
245	MySQL doesn't support views with nested table columns.
296	Amazon SCT can't convert tables that aren't valid.
327	MySQL doesn't support the objects column.
348	MySQL doesn't support global temporary tables.

Constraints



Amazon SCT automatically converts most constraints because Oracle and Amazon Aurora MySQL-Compatible Edition (Aurora MySQL) support the entry level ANSI compliance. These items include primary keys, foreign keys, null constraints, unique constraints, and default constraints with some exceptions. Manual conversions are required for some foreign key cascading options. Amazon SCT replaces check constraints with triggers, and some default expressions for DateTime columns aren't supported for automatic conversion. Amazon SCT can't automatically convert complex expressions for other default values.

For more information, see [Table Constraints](#).

Action code	Action message
202	MySQL doesn't support foreign keys with different types of columns or with referenced columns.
203	Amazon SCT can't convert foreign keys with the SET NULL action for columns that have the NOT NULL constraint.

Action code	Action message
204	Amazon SCT can't convert foreign keys with BLOB and TEXT columns.
220	MySQL doesn't support the record type.
325	Amazon SCT uses triggers to emulate check constraints because MySQL doesn't support them.
326	MySQL doesn't support constraints with the status set to DISABLED.
591 / 593	Amazon SCT can't convert the ROWID usage. This object uses the ROWID column from the %s table.

Data types



Data type syntax and rules are similar between Oracle and Aurora MySQL. Amazon SCT automatically converts most of data type syntax and rules. Date and time handling paradigms are different for Oracle and Aurora MySQL and require manual verification or conversion. Also note that because of differences in data type behavior between Oracle and Aurora MySQL, manual verification and strict testing are highly recommended.

For more information, see [Data Types](#).

Action code	Action message
25	MySQL doesn't support assignment values for variables of the INTERVAL datatype.

Action code	Action message
28	Amazon SCT can't convert the variable declaration of the %s unsupported data type.
29	Amazon SCT can't convert the reference of the %s unsupported data type.
30	Amazon SCT can't convert the usage of the %s unsupported data type.
33	MySQL doesn't support fractional seconds for <code>TIMESTAMP</code> literals.
212	MySQL doesn't support the <code>BFILE</code> data type.

Common table expressions



Aurora MySQL version 5.7 doesn't support common table expressions. Amazon SCT can't automatically convert common table expressions.

For workarounds using traditional SQL syntax, see [Common Table Expressions](#).

Action code	Action message
127	MySQL doesn't support recursive <code>WITH</code> clauses.

Cursors



Amazon SCT automatically converts the most commonly used cursor operations. These operations include forward-only, read only cursors, and the `DECLARE CURSOR`, `CLOSE CURSOR`, and `FETCH NEXT` operations. Modifications through cursors and non-forward-only fetches, which aren't supported by Aurora MySQL, require manual conversions.

For more information, see [Cursors](#).

Action code	Action message
31	Amazon SCT can't convert CURSOR expressions.
84	Amazon SCT doesn't convert the <code>SQL%ISOPE N</code> cursor attribute because this is the default behavior in MySQL.
85	MySQL doesn't support the <code>SQL%BULK_ROWCOUNT</code> cursor attribute.
297	MySQL doesn't support <code>%ROWTYPE</code> attributes.
330	MySQL doesn't support global cursors. Amazon SCT replaces global cursors with local cursors.
337	MySQL doesn't support variables of the <code>SYS_REFCURSOR</code> type.
343	Amazon SCT can't convert SELECT statements for cursors.
354	Amazon SCT can't convert dynamic SQL for the <code>REF_CURSOR</code> variable.
596	Converted code might produce different results compared to the source code. If <code>SQL%ROWCOUNT</code> refers to INSERT or DELETE statements, make sure that you use <code>FOUND_ROWS()</code> instead of <code>ROW_COUNT()</code> .

Action code	Action message
598	MySQL doesn't support RETURN clauses in cursors.

Transaction isolation



Aurora MySQL supports the following four transaction isolation levels specified in the SQL:92 standard: READ UNCOMMITTED, READ COMMITTED, REPEATABLE READ, and SERIALIZABLE. Amazon Schema Conversion Tool (Amazon SCT) automatically converts all these transaction isolation levels. Amazon SCT also converts BEGIN, COMMIT, and ROLLBACK commands that use slightly different syntax. Manual conversion is required for named, marked, and delayed durability transactions that aren't supported by Aurora MySQL.

For more information, see [Transactions](#).

Action code	Action message
235	MySQL doesn't support PRAGMA options.
302	MySQL doesn't support NOWAIT clauses in LOCK TABLE statements.
346	MySQL doesn't support LOCK TABLE statements inside stored procedures.
350	Amazon SCT can't convert statements such as START TRANSACTION , COMMIT, or ROLLBACK.

Stored procedures



Aurora MySQL stored procedures provide very similar functionality to Oracle stored procedures. Amazon SCT automatically converts Oracle stored procedures. Manual conversion is required for procedures that use RETURN values and some less common EXECUTE options such as RECOMPILE and RESULTS SETS.

For more information, see [Stored Procedures](#).

Action code	Action message
27	The package body doesn't include source code.
152	Converted code might not cover all built-in exception names.
234	MySQL doesn't support the EXCEPTION declaration.
253	MySQL doesn't support the %s function with two parameters.
266	MySQL doesn't support the %s function with analytic clauses.
329	MySQL doesn't support RAISE statements.
331	MySQL doesn't support global user-defined exceptions.
333	MySQL doesn't support exception blocks in initialization blocks in packages.
335	MySQL doesn't support GOTO operators.

Action code	Action message
340	MySQL doesn't support the %s function.
342	MySQL doesn't support the %s exception.
345	Converted code might not cover all conditions.
350	Amazon SCT can't convert statements such as <code>START TRANSACTION</code> , <code>COMMIT</code> , or <code>ROLLBACK</code> .
590	Amazon SCT converted the function as procedure.

Triggers



Aurora MySQL supports BEFORE and AFTER triggers for INSERT, UPDATE, and DELETE. Aurora MySQL triggers differ substantially from Oracle triggers. However, for most common use cases, Amazon SCT can migrate triggers with minimal code changes. Although Amazon SCT can automatically migrate trigger code, manual inspection and potential code modifications may be required because Aurora MySQL triggers run once for each row, not once for each statement such as triggers in Oracle.

For more information, see [Triggers](#).

Action code	Action message
236	MySQL doesn't support <code>INSTEAD OF</code> triggers.
237	MySQL doesn't support statement triggers.
238	MySQL doesn't support <code>REFERENCING</code> clauses.

Action code	Action message
239	MySQL doesn't support triggers with WHEN conditions.
240	MySQL doesn't support triggers on nested table columns in views.
241	MySQL doesn't support triggers with FOLLOWS and PRECEDES clauses.
242	MySQL doesn't support compound triggers.
243	MySQL doesn't support UPDATE OF clauses.
244	MySQL doesn't support conditional predicates.
306	Amazon SCT can't convert a trigger that isn't valid.
310	MySQL doesn't support triggers for views.
311	MySQL doesn't support system triggers.
312	MySQL doesn't support DISABLED clauses.
313	MySQL doesn't support action-type clauses in triggers.
314	MySQL doesn't support cross edition triggers.
316	MySQL doesn't support the apply-server-only property.
317	MySQL doesn't support PARENT referencing clauses.
415	MySQL doesn't support system triggers.
524	MySQL doesn't support triggers for multiple events.

Action code	Action message
588	MySQL doesn't support multiple triggers for a single event. Amazon SCT merged triggers in one trigger.

Sequences



Although the syntax for Oracle IDENTITY and Aurora MySQL AUTO_INCREMENT auto-enumeration columns differs significantly, Amazon SCT can automatically convert it. Some limitations imposed by Aurora MySQL require manual conversion such as explicit SEED and INCREMENT auto-enumeration columns that aren't part of the primary key and the table-independent SEQUENCE objects.

For more information, see [Oracle Sequences and Identity Columns and MySQL Sequences and AUTO INCREMENT Columns](#).

Action code	Action message
341	MySQL doesn't support sequences.

Date and time functions



Amazon SCT automatically converts the most commonly used date and time functions despite the significant difference in syntax. Be aware of differences in data types, time zone awareness, and locale handling as well the functions themselves, and inspect the expression value output carefully. Some less commonly used options such as millisecond, nanosecond, and time zone offsets require manual conversion.

Action code	Action message
213	Amazon SCT expanded fractional seconds support for TIME, DATETIME, and TIMESTAMP values with up to 6 digits of precision.
214	MySQL doesn't support data types that store time zone information. The DATETIME data type stores timestamps in the MySQL server time zone.
216	Amazon SCT expanded fractional seconds support for TIME, DATETIME, and TIMESTAMP values with up to 6 digits of precision. MySQL doesn't support data types that store time zone information.

User-defined types



Aurora MySQL 5.7 doesn't support user defined types and user-defined table-valued parameters. Amazon SCT can convert standard user defined types by replacing it with their base types, but manual conversion is required for user defined table types, which are used for table valued parameters for stored procedures.

For more information, see [User-Defined Types](#).

Action code	Action message
196	MySQL doesn't support object tables.
218	MySQL doesn't support user types.

Synonyms



Aurora MySQL version 5.7 doesn't support synonyms. Amazon SCT can't automatically convert synonyms.

Action code	Action message
352	MySQL doesn't support synonyms. Amazon SCT replaces synonyms with fully-qualified object names.

XML



Aurora MySQL provides minimal support for XML, but it does offer a native JSON data type and more than 25 dedicated JSON functions. Despite these differences, the most commonly used basic XML functions can be automatically migrated by Amazon SCT. Some options such as EXPLICIT, used in functions or with subqueries, require manual conversion.

For more information, see [XML](#).

Action code	Action message
194	MySQL doesn't support XMLTYPE tables.
195	MySQL doesn't support the XMLDATA pseudocolumn.
303	MySQL doesn't support the XMLTable function.

MERGE



Aurora MySQL version 5.7 doesn't support the MERGE statement. Amazon SCT can't automatically convert MERGE statements. Manual conversion is straightforward in most cases.

Action code	Action message
102	MySQL doesn't support MERGE statements.

Query hints



Amazon SCT can automatically convert basic query hints such as index hints, except for DML statements. Note that specific optimizations used for Oracle may be completely inapplicable to a new query optimizer. It is recommended to start migration testing with all hints removed. Then, selectively apply hints as a last resort if other means such as schema, index, and query optimizations have failed. Plan guides aren't supported by Aurora MySQL.

For more information, see [Database Hints](#).

Action code	Action message
103	Amazon SCT can't convert hints. MySQL doesn't support the %s hint.

Indexes



Amazon SCT automatically converts basic non-clustered indexes, which are the most commonly used type of indexes. User-defined clustered indexes aren't supported by Aurora MySQL because they are always created for the primary key. In addition, filtered indexes, indexes with included columns, and some Oracle specific index options can't be migrated automatically and require manual conversion.

For more information, see [Indexes](#).

Action code	Action message
205	MySQL has reached the limit of the internal InnoDB maximum key length.
206	MySQL doesn't support bitmap indexes.
207	MySQL doesn't support function indexes.
208	MySQL doesn't support domain indexes.
328	Amazon SCT can't convert indexes that aren't valid.

Partitioning



Because Aurora MySQL stores each table in its own file, and because file management is performed by Amazon and can't be modified, some of the physical aspects of partitioning in Oracle don't apply to Aurora MySQL. Because of the vast differences between partition creation, query, and management between Aurora MySQL and Oracle, Amazon SCT doesn't automatically convert table and index partitions. These items require manual conversion.

For more information, see [Table Partitioning](#).

Action code	Action message
201	MySQL doesn't support partition types that are implemented in your source code.
699	MySQL doesn't support not allowed partitions functions.

Materialized views



Aurora MySQL 5.7 doesn't support materialized views. Amazon SCT can't automatically convert materialized views.

For more information, see [Oracle Materialized Views and MySQL Summary Tables or Views](#).

Action code	Action message
94	MySQL doesn't support materialized views.
95	MySQL doesn't support the usage of materialized views.

Views



Although the basic syntax for creating a view in Oracle and Aurora MySQL is almost identical, there are some sub-options that can differ significantly, requiring additional manual migration tasks.

For more information, see [Views](#).

Action code	Action message
75	MySQL doesn't support read-only views.
93	MySQL doesn't support UPDATE statements for views.
97	MySQL doesn't support DELETE statements for views.
320	Amazon SCT can't convert a view that isn't valid.
321	MySQL doesn't support object views.
323	MySQL doesn't support subviews under a superview.
324	MySQL doesn't support editioning views.
583	MySQL doesn't support constraints for views.

UTL_Mail and UTL_SMTP



Aurora MySQL doesn't provide native support for sending emails from the database.

For more information, see [Database Mail](#).

Action code	Action message
81	MySQL doesn't support sending SMS notifications.
82	MySQL doesn't support sending emails.

Database Links



Aurora MySQL doesn't support remote data access. Connectivity between schemas is trivial, but connectivity to other instances requires a custom solution. Amazon SCT can't automatically convert database links.

For more information, see [Database Links](#).

Action code	Action message
600	MySQL doesn't support the usage of database links.

PLSQL



Amazon SCT automatically converts the most commonly used SQL statements because Oracle and Aurora MySQL support the entry level ANSI compliance. Some changes may be required for DML related to ERROR LOG, subquery, and partitions.

Action code	Action message
63	Amazon SCT can't convert UPDATE statements with multiple-column subqueries in SET clauses.
64	MySQL doesn't support UPDATE statements with ERROR LOG clauses.
65	MySQL doesn't support UPDATE statements for subqueries.

Action code	Action message
66	MySQL doesn't support UPDATE statements for RETURNING INTO clauses.
67	MySQL doesn't support DELETE statements with ERROR LOG clauses.
68	MySQL doesn't support DELETE statements for subqueries.
69	MySQL doesn't support DELETE statements for RETURNING INTO clauses.
70	MySQL doesn't support INSERT statements with ERROR LOG clauses.
71	MySQL doesn't support INSERT statements for subqueries.
72	MySQL doesn't support INSERT statements for RETURNING INTO clauses.
77	MySQL doesn't support PIVOT clauses for SELECT statements.
78	MySQL doesn't support UNPIVOT clauses for SELECT statements.
87	MySQL doesn't support RETURNING BULK COLLECT INTO clauses.
89	MySQL doesn't support INSERT statements for views.
90	MySQL doesn't support INSERT statements for subpartitions.
122	MySQL doesn't support hierarchical queries.

Action code	Action message
125	MySQL doesn't support GROUPING SETS statements.
128	MySQL doesn't support ORACLE FLASHBACK VERSION QUERY.
138	MySQL doesn't support FOR UPDATE OF clauses.
139	MySQL doesn't support FOR UPDATE SKIP LOCKED clauses.
140	MySQL doesn't support BULK COLLECT INTO clauses.
141	MySQL doesn't support ORDER BY ... NULLS FIRST clauses.
143	MySQL doesn't support FOR UPDATE NOWAIT clauses.
144	MySQL doesn't support FOR UPDATE WAIT clauses.
585	Amazon SCT can't convert outer join inside a correlated query.
594	MySQL doesn't support LATERAL, CROSS APPLY, and OUTER APPLY correlated inline views.
599	MySQL doesn't support CURRENT OF clauses for data manipulation language queries that are in the body of a cursor loop.

EXECUTE IMMEDIATE



There is a major difference between Oracle and Aurora MySQL for the EXECUTE IMMEDIATE statement. In MySQL, this statement must be used after a PREPARE command. Running SQL with results and bind variables, and running anonymous blocks aren't supported.

For more information, see [Execute Immediate](#).

Action code	Action message
88	MySQL doesn't support EXECUTE IMMEDIATE statements with BULK COLLECT.
334	MySQL doesn't support EXECUTE IMMEDIATE dynamic SQL statements.
336	MySQL doesn't support EXECUTE IMMEDIATE dynamic SQL statements with the %s clause.

DBMS_OUTPUT



Aurora MySQL doesn't provide native support for the dbms_output procedure. Use the RAISE command instead.

For more information, see [DBMS_OUTPUT](#).

Action code	Action message
332	MySQL doesn't support the dbms_output.put_line procedure.

Action code	Action message
349	MySQL doesn't support the <code>dbms_output.put</code> procedure.

Amazon Database Migration Service

The Amazon Database Migration Service (Amazon DMS) helps you migrate databases to Amazon quickly and securely. The source database remains fully operational during the migration, minimizing downtime to applications that rely on the database. The Amazon Database Migration Service can migrate your data to and from most widely-used commercial and open-source databases.

The service supports homogenous migrations such as Oracle to Amazon RDS for Oracle as well as heterogeneous migrations between different database platforms such as Oracle to Amazon Aurora MySQL. You can also use Amazon DMS to stream data to Amazon Redshift, Amazon DynamoDB, and Amazon S3 from any of the supported sources, which are Amazon Aurora, PostgreSQL, MySQL, MariaDB, Oracle Database, SAP ASE, SQL Server, IBM DB2 LUW, and MongoDB, enabling consolidation and easy analysis of data in a petabyte-scale data warehouse. The Amazon Database Migration Service can also be used for continuous data replication with high availability.

For Amazon DMS pricing, see [Database Migration Service pricing](#).

For all supported sources for Amazon DMS, see [Sources for data migration](#).

For all supported targets for Amazon DMS, see [Targets for data migration](#).

Migration tasks performed by Amazon DMS

In a traditional solution, you need to perform capacity analysis, procure hardware and software, install and administer systems, and test and debug the installation. Amazon DMS automatically manages the deployment, management, and monitoring of all hardware and software needed for your migration. You can start your migration within minutes of starting the Amazon DMS configuration process.

With Amazon DMS, you can scale up (or scale down) your migration resources as needed to match your actual workload. For example, if you determine that you need additional storage, you can easily increase your allocated storage and restart your migration, usually within minutes. On the

other hand, if you discover that you aren't using all of the resource capacity you configured, you can easily downsize to meet your actual workload.

Amazon DMS uses a pay-as-you-go model. You only pay for Amazon DMS resources while you use them as opposed to traditional licensing models with up-front purchase costs and ongoing maintenance charges.

Amazon DMS automatically manages all of the infrastructure that supports your migration server including hardware and software, software patching, and error reporting.

Amazon DMS provides automatic failover. If your primary replication server fails for any reason, a backup replication server can take over with little or no interruption of service.

Amazon DMS can help you switch to a modern, perhaps more cost-effective database engine than the one you are running now. For example, Amazon DMS can help you take advantage of the managed database services provided by Amazon RDS or Amazon Aurora. Or, it can help you move to the managed data warehouse service provided by Amazon Redshift, NoSQL platforms like Amazon DynamoDB, or low-cost storage platforms like Amazon S3. Conversely, if you want to migrate away from old infrastructure but continue to use the same database engine, Amazon DMS also supports that process.

Amazon DMS supports nearly all of modern popular DBMS engines as data sources, including Oracle, Microsoft SQL Server, MySQL, MariaDB, PostgreSQL, Db2 LUW, SAP, MongoDB, and Amazon Aurora.

Amazon DMS provides a broad coverage of available target engines including Oracle, Microsoft SQL Server, PostgreSQL, MySQL, Amazon Redshift, SAP ASE, Amazon S3, and Amazon DynamoDB.

You can migrate from any of the supported data sources to any of the supported data targets. Amazon DMS supports fully heterogeneous data migrations between the supported engines.


Amazon DMS ensures that your data migration is secure. Data at rest is encrypted with Amazon Key Management Service (Amazon KMS) encryption. During migration, you can use Secure Socket Layers (SSL) to encrypt your in-flight data as it travels from source to target.

How Amazon DMS works

At its most basic level, Amazon DMS is a server in the Amazon Cloud that runs replication software. You create a source and target connection to tell Amazon DMS where to extract from and load to. Then, you schedule a task that runs on this server to move your data. Amazon DMS creates the tables and associated primary keys if they don't exist on the target. You can pre-create the target

tables manually if you prefer. Or you can use Amazon SCT to create some or all of the target tables, indexes, views, triggers, and so on.

The following diagram illustrates the Amazon DMS process.



For more information about Amazon DMS, see [What is Database Migration Service?](#) and [Best practices for Database Migration Service](#).

Amazon RDS on Outposts

Note

This topic is related to Amazon Relational Database Service (Amazon RDS) and isn't supported with Amazon Aurora.


Amazon RDS on Outposts is a fully managed service that offers the same Amazon infrastructure, Amazon services, APIs, and tools to virtually any data center, co-location space, or on-premises facility for a truly consistent hybrid experience. Amazon RDS on Outposts is ideal for workloads that require low latency access to on-premises systems, local data processing, data residency, and migration of applications with local system inter-dependencies.

When you deploy Amazon RDS on Outposts, you can run Amazon RDS on premises for low latency workloads that need to be run in close proximity to your on-premises data and applications. Amazon RDS on Outposts also enables automatic backup to an Amazon Region. You can manage

Amazon RDS databases both in the cloud and on premises using the same Amazon Management Console, APIs, and CLI. Amazon RDS on Outposts supports Microsoft SQL Server, MySQL, and PostgreSQL database engines, with support for additional database engines coming soon.

How it works

Amazon RDS on Outposts enables you to run Amazon RDS in your on-premises or co-location site. You can deploy and scale an Amazon RDS database instance in Outposts just as you do in the cloud, using the Amazon Management Console, APIs, or CLI. Amazon RDS databases in Outposts are encrypted at rest using Amazon KMS keys. Amazon RDS automatically stores all automatic backups and manual snapshots in the Amazon Region.



This option is helpful when you need to run Amazon RDS on premises for low latency workloads that need to be run in close proximity to your on-premises data and applications.

For more information, see [Amazon Outposts Family](#), [Amazon RDS on Outposts](#), and [Create Amazon RDS DB Instances on Outposts](#).

Amazon RDS Proxy

Amazon RDS Proxy is a fully managed, highly available database proxy for Amazon Relational Database Service (RDS) that makes applications more scalable, more resilient to database failures, and more secure.

Many applications, including those built on modern server-less architectures, can have many open connections to the database server, and may open and close database connections at a high rate, exhausting database memory and compute resources. Amazon RDS Proxy allows applications to pool and share connections established with the database, improving database efficiency and application scalability. With Amazon RDS Proxy, fail-over times for Aurora and Amazon RDS databases are reduced by up to 66%. You can manage database credentials, authentication, and access through integration with Amazon Secrets Manager and Amazon Identity and Access Management (IAM).

You can turn on Amazon RDS Proxy for most applications with no code changes. You don't need to provision or manage any additional infrastructure. Pricing is simple and predictable: you pay for each vCPU of the database instance for which the proxy is enabled. Amazon RDS Proxy is now generally available for Aurora MySQL, Aurora PostgreSQL, Amazon RDS for MySQL, and Amazon RDS for PostgreSQL.

Amazon RDS Proxy benefits

Amazon RDS Proxy is a fully managed service that provides the following benefits:

- **Improved application performance** — Amazon RDS Proxy manages a connection pooling which helps with reducing the stress on database compute and memory resources that typically occurs when new connections are established and it is useful to efficiently support a large number and frequency of application connections.
- **Increased application availability** — By automatically connecting to a new database instance while preserving application connections Amazon RDS Proxy can reduce fail-over time by 66%.
- **Manageable application security** — Amazon RDS Proxy also enables you to centrally manage database credentials using Amazon Secrets Manager.
- **Fully managed** — Amazon RDS Proxy gives you the benefits of a database proxy without requiring additional burden of patching and managing your own proxy server.
- **Fully compatible with your database** — Amazon RDS Proxy is fully compatible with the protocols of supported database engines, so you can deploy Amazon RDS Proxy for your application without making changes to your application code.
- **Available and durable** — Amazon RDS Proxy is highly available and deployed over multiple Availability Zones (AZs) to protect you from infrastructure failure.

How Amazon RDS Proxy works

The following diagram shows how Amazon RDS Proxy works.



For more information, see [Amazon RDS Proxy for Scalable Serverless Applications](#) and [Amazon RDS Proxy](#).

Amazon Aurora Serverless v1

Amazon Aurora Serverless version 1 (v1) is an on-demand autoscaling configuration for Amazon Aurora. An Aurora Serverless DB cluster is a DB cluster that scales compute capacity up and down based on your application's needs. This contrasts with Aurora provisioned DB clusters, for which you manually manage capacity. Aurora Serverless v1 provides a relatively simple, cost-effective option for infrequent, intermittent, or unpredictable workloads. It is cost-effective because it automatically starts up, scales compute capacity to match your application's usage, and shuts down when it's not in use.

To learn more about pricing, see Serverless Pricing under MySQL-Compatible Edition or PostgreSQL-Compatible Edition on the [Amazon Aurora pricing page](#).

Aurora Serverless v1 clusters have the same kind of high-capacity, distributed, and highly available storage volume that is used by provisioned DB clusters. The cluster volume for an Aurora Serverless v1 cluster is always encrypted. You can choose the encryption key, but you can't turn off encryption. That means that you can perform the same operations on an Aurora Serverless v1 that you can on encrypted snapshots.

Aurora Serverless v1 provides the following advantages:

- **Simpler than provisioned** — Aurora Serverless v1 removes much of the complexity of managing DB instances and capacity.


- **Scalable** — Aurora Serverless v1 seamlessly scales compute and memory capacity as needed, with no disruption to client connections.
- **Cost-effective** — When you use Aurora Serverless v1, you pay only for the database resources that you consume, on a per-second basis.
- **Highly available storage** — Aurora Serverless v1 uses the same fault-tolerant, distributed storage system with six-way replication as Aurora to protect against data loss.

Aurora Serverless v1 is designed for the following use cases:

- **Infrequently used applications** — You have an application that is only used for a few minutes several times for each day or week, such as a low-volume blog site. With Aurora Serverless v1, you pay for only the database resources that you consume on a per-second basis.
- **New applications** — You're deploying a new application and you're unsure about the instance size you need. By using Aurora Serverless v1, you can create a database endpoint and have the database automatically scale to the capacity requirements of your application.
- **Variable workloads** — You're running a lightly used application, with peaks of 30 minutes to several hours a few times each day, or several times for each year. Examples are applications for human resources, budgeting, and operational reporting applications. With Aurora Serverless v1, you no longer need to provision for peak or average capacity.
- **Unpredictable workloads** — You're running daily workloads that have sudden and unpredictable increases in activity. An example is a traffic site that sees a surge of activity when it starts raining. With Aurora Serverless v1, your database automatically scales capacity to meet the needs of the application's peak load and scales back down when the surge of activity is over.
- **Development and test databases** — Your developers use databases during work hours but don't need them on nights or weekends. With Aurora Serverless v1, your database automatically shuts down when it's not in use.
- **Multi-tenant applications** — With Aurora Serverless v1, you don't have to individually manage database capacity for each application in your fleet. Aurora Serverless v1 manages individual database capacity for you.

This process takes almost no time. Because the storage is shared between nodes Aurora can scale up or down in seconds for most workloads. The service currently has autoscaling thresholds of 1.5 minutes to scale up and 5 minutes to scale down. That means metrics must exceed the limits for 1.5 minutes to trigger a scale up or fall below the limits for 5 minutes to trigger a scale down. The cool-down period between scaling activities is 5 minutes to scale up and 15 minutes to scale

down. Before scaling can happen the service has to find a “scaling point” which may take longer than anticipated if you have long-running transactions. Scaling operations are transparent to the connected clients and applications since existing connections and session state are transferred to the new nodes. The only difference with pausing and resuming is a higher latency for the first connection, typically around 25 seconds. You can find more details in the documentation.



Amazon Aurora Serverless v2

Amazon Aurora Serverless v2 has been architected from the ground up to support serverless DB clusters that are instantly scalable. The Aurora Serverless v2 architecture rests on a lightweight foundation that’s engineered to provide the security and isolation needed in multitenant serverless cloud environments. This foundation has very little overhead so it can respond quickly. It’s also powerful enough to meet dramatic increases in processing demand.

When you create your Aurora Serverless v2 DB cluster, you define its capacity as a range between minimum and maximum number of Aurora capacity units (ACUs):

- **Minimum Aurora capacity units** — The smallest number of ACUs down to which your Aurora Serverless v2 DB cluster can scale.
- **Maximum Aurora capacity units** — The largest number of ACUs up to which your Aurora Serverless v2 DB cluster can scale.

Each ACU provides 2 GiB (gibibytes) of memory (RAM) and associated virtual processor (vCPU) with networking.

Unlike Aurora Serverless v1, which scales by doubling ACUs each time the DB cluster reaches a threshold, Aurora Serverless v2 can increase ACUs incrementally. When your workload demand begins to reach the current resource capacity, your Aurora Serverless v2 DB cluster scales the number of ACUs. Your cluster scales ACUs in the increments required to provide the best performance for the resources consumed.

How to provision

Log in to your [Management Console](#), choose **Amazon RDS**, and then choose **Create database**.

On **Engine options**, for **Engine versions**, choose **Show versions that support Serverless v2**.

Engine options

Engine type [Info](#)

<input checked="" type="radio"/> Amazon Aurora 	<input type="radio"/> MySQL 	<input type="radio"/> MariaDB 
<input type="radio"/> PostgreSQL 	<input type="radio"/> Oracle 	<input type="radio"/> Microsoft SQL Server 

Edition

- Amazon Aurora MySQL-Compatible Edition
- Amazon Aurora PostgreSQL-Compatible Edition

► Replication features [Info](#)

Single-master replication is currently selected

Engine version [Info](#)

View the engine versions that support the following database features.

► Show filters

Available versions (1/60) [Info](#)

Aurora MySQL 3.02.0 (compatible with MySQL 8.0.23) ▼

Choose the capacity settings for your use case.

For more information, see [Amazon Aurora Serverless](#), [Aurora Serverless MySQL Generally Available](#), and [Amazon Aurora PostgreSQL Serverless Now Generally Available](#).

Amazon Aurora Parallel Query

Amazon Aurora Parallel Query is a feature of the Amazon Aurora database that provides faster analytical queries over your current data, without having to copy the data into a separate system. It can speed up queries by up to two orders of magnitude, while maintaining high throughput for your core transactional workload.

While some databases can parallelize query processing across CPUs in one or a handful of servers, Parallel Query takes advantage of Aurora unique architecture to push down and parallelize query processing across thousands of CPUs in the Aurora storage layer. By offloading analytical query processing to the Aurora storage layer, Parallel Query reduces network, CPU, and buffer pool contention with the transactional workload.

Features

Accelerate your analytical queries

In a traditional database, running analytical queries directly on the database means accepting slower query performance and risking a slowdown of your transactional workload, even when running light queries. Queries can run for several minutes to hours, depending on the size of the tables and database server instances. Queries are also slowed down by network latency, since the storage layer may have to transfer entire tables to the database server for processing.

With Amazon Aurora Parallel Query, query processing is pushed down to the Aurora storage layer. The query gains a large amount of computing power, and it needs to transfer far less data over the network. In the meantime, the Amazon Aurora database instance can continue serving transactions with much less interruption. This way, you can run transactional and analytical workloads alongside each other in the same Aurora database, while maintaining high performance.

Query on fresh data

Many analytical workloads require both fresh data and good query performance. For example, operational systems such as network monitoring, cyber-security or fraud detection rely on fresh, real-time data from a transactional database, and can't wait for it to be extracted to a analytics system.

By running your queries in the same database that you use for transaction processing, without degrading transaction performance, Amazon Aurora Parallel Query enables smarter operational decisions with no additional software and no changes to your queries.

Benefits of Using Parallel Query

AURLong Parallel Query feature provides the following benefits:

- Improved I/O performance, due to parallelizing physical read requests across multiple storage nodes.
- Reduced network traffic. Amazon Aurora doesn't transmit entire data pages from storage nodes to the head node and then filter out unnecessary rows and columns afterward. Instead, Aurora transmits compact tuples containing only the column values needed for the result set.
- Reduced CPU usage on the head node, due to pushing down function processing, row filtering, and column projection for the WHERE clause.
- Reduced memory pressure on the buffer pool. The pages processed by Parallel Query aren't added to the buffer pool. This approach reduces the chance of a data-intensive scan evicting frequently used data from the buffer pool.
- Potentially reduced data duplication in your extract, transform, and load (ETL) pipeline, by making it practical to perform long-running analytic queries on existing data.

Important notes

Consider the following when using the AURLong Parallel Query feature:

- **Table formats** — The table row format must be COMPACT; partitioned tables aren't supported.
- **Data types** — The TEXT, BLOB, and GEOMETRY data types aren't supported.
- **DDL** — The table can't have any pending fast online DDL operations.
- **Cost** — You can make use of Parallel Query at no extra charge. However, because it makes direct access to storage, there is a possibility that your IO cost will increase.

For more information, see [Amazon Aurora Parallel Query](#).

Amazon Aurora Backtrack

We've all been there, you need to make a quick, seemingly simple fix to an important production database. You compose the query, give it a once-over, and let it run. Seconds later you realize that you forgot the WHERE clause, dropped the wrong table, or made another serious mistake, and

interrupt the query, but the damage has been done. You take a deep breath, whistle through your teeth, wish that reality came with an Undo option.

Backtracking rewinds the DB cluster to the time you specify. Backtracking isn't a replacement for backing up your DB cluster so that you can restore it to a point in time. However, backtracking provides the following advantages over traditional backup and restore:

- You can easily undo mistakes. If you mistakenly perform a destructive action, such as a DELETE without a WHERE clause, you can backtrack the DB cluster to a time before the destructive action with minimal interruption of service.
- You can backtrack a DB cluster quickly. Restoring a DB cluster to a point in time launches a new DB cluster and restores it from backup data or a DB cluster snapshot, which can take hours. Backtracking a DB cluster doesn't require a new DB cluster and rewinds the DB cluster in minutes.
- You can explore earlier data changes. You can repeatedly backtrack a DB cluster back and forth in time to help determine when a particular data change occurred. For example, you can backtrack a DB cluster three hours and then backtrack forward in time one hour. In this case, the backtrack time is two hours before the original time.

Amazon Aurora uses a distributed, log-structured storage system (read [Design Considerations for High Throughput Cloud-Native Relational Databases](#) to learn a lot more); each change to your database generates a new log record, identified by a Log Sequence Number (LSN). Enabling the backtrack feature provisions a FIFO buffer in the cluster for storage of LSNs. This allows for quick access and recovery times measured in seconds.

When you create a new Aurora MySQL DB cluster, backtracking is configured when you choose **Enable Backtrack** and specify a **Target Backtrack window** value that is greater than zero in the Backtrack section.

To create a DB cluster, follow the instructions in [Creating an Amazon Aurora DB cluster](#). The following image shows the Backtrack section.

Backtrack

Backtrack lets you quickly rewind the DB cluster to a specific point in time, without having to create another DB cluster. [Info](#)

Enable Backtrack

Enabling Backtrack will charge you for storing the changes you make for backtracking.

Target Backtrack window

The Backtrack window determines how far back in time you could go. Aurora will try to retain enough log information to support that window of time. [Info](#)

 hours (up to 72)

Typical user cost

The cost of Backtrack depends on how often you are updating your database. This is an estimate based on typical workloads for your selected instance size (db.r5.large). [Info](#)

\$ 14.38 USD / month

After a production error, you can simply pause your application, open up the Aurora console, select the cluster, and choose **Backtrack DB cluster**.

Then you select **Backtrack** and choose the point in time just before your epic fail, and choose **Backtrack DB cluster**.

Backtrack DB cluster

Rewinds the DB cluster to a previous point in time without creating a new DB cluster.

Earliest restorable time is June 16, 2021 at 8:53:02 PM UTC-4 (Local) ⓘ

Date	Time				
<input type="text" value="June 16, 2021"/>	<input type="text" value="20"/> ▼	:	<input type="text" value="53"/> ▼	:	<input type="text" value="02"/> ▼

UTC-4

The next available time will be used if the specified time is not available.

 Your DB cluster is unavailable during the Backtrack process, which typically takes a few minutes.

Then you wait for the rewind to take place, unpause your application and proceed as if nothing had happened. When you initiate a backtrack, Aurora will pause the database, close any open connections, drop uncommitted writes, and wait for the backtrack to complete. Then it will resume

normal operation and be able to accept requests. The instance state will be backtracking while the rewind is underway.

Backtrack window

With backtracking, there is a target backtrack window and an actual backtrack window:

- The target backtrack window is the amount of time you want to be able to backtrack your DB cluster. When you enable backtracking, you specify a target backtrack window. For example, you might specify a target backtrack window of 24 hours if you want to be able to backtrack the DB cluster one day.
- The actual backtrack window is the actual amount of time you can backtrack your DB cluster, which can be smaller than the target backtrack window. The actual backtrack window is based on your workload and the storage available for storing information about database changes, called change records.

As you make updates to your Aurora DB cluster with backtracking enabled, you generate change records. Aurora retains change records for the target backtrack window, and you pay an hourly rate for storing them. Both the target backtrack window and the workload on your DB cluster determine the number of change records you store. The workload is the number of changes you make to your DB cluster in a given amount of time. If your workload is heavy, you store more change records in your backtrack window than you do if your workload is light.

You can think of your target backtrack window as the goal for the maximum amount of time you want to be able to backtrack your DB cluster. In most cases, you can backtrack the maximum amount of time that you specified. However, in some cases, the DB cluster can't store enough change records to backtrack the maximum amount of time, and your actual backtrack window is smaller than your target. Typically, the actual backtrack window is smaller than the target when you have extremely heavy workload on your DB cluster. When your actual backtrack window is smaller than your target, we send you a notification.

When backtracking is turned on for a DB cluster, and you delete a table stored in the DB cluster, Aurora keeps that table in the backtrack change records. It does this so that you can revert back to a time before you deleted the table. If you don't have enough space in your backtrack window to store the table, the table might be removed from the backtrack change records eventually.

Backtracking limitations

The following limitations apply to backtracking:

- Backtracking an Aurora DB cluster is available in certain Amazon Regions and for specific Aurora MySQL versions only. For more information, see [Backtracking in Aurora](#).
- Backtracking is only available for DB clusters that were created with the Backtrack feature enabled. You can enable the Backtrack feature when you create a new DB cluster or restore a snapshot of a DB cluster. For DB clusters that were created with the Backtrack feature enabled, you can create a clone DB cluster with the Backtrack feature enabled. Currently, you can't perform backtracking on DB clusters that were created with the Backtrack feature turned off.
- The limit for a backtrack window is 72 hours.
- Backtracking affects the entire DB cluster. For example, you can't selectively backtrack a single table or a single data update.
- Backtracking isn't supported with binary log (binlog) replication. Cross-Region replication must be turned off before you can configure or use backtracking.
- You can't backtrack a database clone to a time before that database clone was created. However, you can use the original database to backtrack to a time before the clone was created. For more information about database cloning, see [Cloning an Aurora DB cluster volume](#).
- Backtracking causes a brief DB instance disruption. You must stop or pause your applications before starting a backtrack operation to ensure that there are no new read or write requests. During the backtrack operation, Aurora pauses the database, closes any open connections, and drops any uncommitted reads and writes. It then waits for the backtrack operation to complete.
- Backtracking isn't supported for the following Amazon Regions:
 - Africa (Cape Town)
 - China (Ningxia)
 - Asia Pacific (Hong Kong)
 - Europe (Milan)
 - Europe (Stockholm)
 - Middle East (Bahrain)
 - South America (São Paulo)
- You can't restore a cross-region snapshot of a backtrack-enabled cluster in an Amazon Region that doesn't support backtracking.
- You can't use backtrack with Aurora multi-master clusters.

- If you perform an in-place upgrade for a backtrack-enabled cluster from Aurora MySQL version 1 to version 2, you can't backtrack to a point in time before the upgrade happened.

For more information, see: [Amazon Aurora Backtrack — Turn Back Time](#).

SQL and PL/SQL

This section provides reference pages to Oracle and MySQL functions, statements, and other commands.

Topics

- [Single-row and aggregate Oracle and MySQL functions](#)
- [Oracle and MySQL CREATE TABLE AS SELECT statement](#)
- [Oracle and MySQL Common Table Expressions](#)
- [Oracle sequences and identity columns and MySQL sequences and AUTO INCREMENT columns](#)
- [Oracle and MySQL INSERT FROM SELECT statement](#)
- [Multi-Version Concurrency Control](#)
- [Oracle MERGE statement and MySQL equivalent](#)
- [Oracle OLAP functions and MySQL Window functions](#)
- [Oracle Transaction Model and MySQL Transactions](#)
- [Oracle anonymous block and MySQL transactions or procedures](#)
- [Conversion functions](#)
- [Oracle and MySQL cursors](#)
- [Oracle DBMS_DATAPUMP and MySQL integration with Amazon S3](#)
- [Oracle DBMS_OUTPUT and MySQL SELECT](#)
- [Oracle DBMS_RANDOM and MySQL RAND function](#)
- [Oracle DBMS_REDEFINITION and MySQL tables and triggers](#)
- [Oracle DBMS_SQL](#)
- [Oracle EXECUTE IMMEDIATE and MySQL EXECUTE and PREPARE statements](#)
- [Oracle procedures and functions and MySQL stored procedures](#)
- [Oracle and MySQL regular expressions](#)
- [Oracle TIMEZONE data type and functions and MySQL CONVERT_TZ function](#)
- [Oracle and MySQL user-defined functions](#)
- [Oracle UTL_FILE and MySQL integration with Amazon S3](#)
- [Oracle UTL_MAIL or UTL_SMTP and Amazon Simple Notification Service](#)

Single-row and aggregate Oracle and MySQL functions

Single-row and aggregate functions are essential SQL constructs that perform operations on individual rows or groups of rows, respectively. The following sections compare Oracle and MySQL single-row and aggregate functions.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	MySQL doesn't support all functions. These unsupported functions require manual creation.

Oracle usage

Oracle provides two main categories of built-in SQL functions based on the number of rows used as input and generated as output.

- **Single-row** or scalar functions return a single result for each row of the queried table or view. You can use them with a SELECT statement in the WHERE clause, the START WITH clause, the CONNECT BY clause, and the HAVING clause. The single-row functions are divided into groups according to data types such as NUMERIC functions, CHAR functions, and DATETIME functions.
- **Aggregative** or group functions are used to summarize a group of values into a single result. Examples include AVG, MIN, MAX, SUM, COUNT, LISTAGG, FIRST, and LAST.

See the following section for a comparison of Oracle and MySQL single-row functions.

Oracle 19 adds ability to eliminate duplicate items in LISTAGG function results with new DISTINCT keyword.

Oracle 19 introduces several new bitmap SQL aggregate functions such as BITMAP_BUCKET_NUMBER, BITMAP_BIT_POSITION and BITMAP_CONSTRUCT_AGG. These functions help speed up COUNT DISTINCT operations.

For more information, see [Single-Row Functions](#) and [Aggregate Functions](#) in *Oracle documentation*.

MySQL usage

MySQL provides an extensive list of single-row and aggregation functions. Some are similar to their Oracle counterparts by name and functionality, or under a different name but with similar functionality. Other functions can have identical names to their Oracle counterparts, but exhibit different functionality. In the following tables, the Equivalent column indicates functional equivalency.

Numeric functions

Oracle function and definition	MySQL function and definition	Equivalent
ABS — Absolute value of n: abs (-11.3) = 11.3.	ABS — Absolute value of n: abs (-11.3) = 11.3.	Yes
CEIL — Returns the smallest integer that is greater than or equal to n: ceil (-24.9) = -24.	CEIL — Returns the smallest integer that is greater than or equal to n: ceil (-24.9) = -24.	Yes
FLOOR — Returns the largest integer equal to or less than n: floor (-43.7) = -44.	FLOOR — Returns the largest integer equal to or less than n: floor (-43.7) = -44.	Yes
MOD — Remainder of n2 divided by n1: mod(10, 3) = 1.	MOD — Remainder of n2 divided by n1: mod(10, 3) = 1.	Yes
ROUND — Returns n rounded to integer places to the right of the decimal point: round (3.49, 1) = 3.5.	ROUND — Returns n rounded to integer places to the right of the decimal point: round (3.49, 1) = 3.5.	Yes
TRUNC — Returns n1 truncated to n2 decimal	TRUNCATE — Returns n1 truncated to n2 decimal	Yes

Oracle function and definition	MySQL function and definition	Equivalent
places: <code>trunc(13.5) = 13.</code>	places: <code>trunc(13.5) = 13.</code>	

Character functions

Oracle function and definition	MySQL function and definition	Equivalent
CONCAT — Returns char1 concatenated with char2: <code>concat('a', 1) # a1 .</code>	CONCAT — Returns char1 concatenated with char2: <code>concat('a', 1) # a1 .</code>	Yes
LOWER and UPPER — Returns char, with all letters lowercase or uppercase: <code>lower ('MR. Smith') # mr. smith.</code>	LOWER and UPPER — Returns char, with all letters lowercase or uppercase: <code>lower ('MR. Smith') # mr. smith.</code>	Yes
LPAD and RPAD — Returns expr1, left or right padded to length n characters with the sequence of characters in expr2: <code>LPAD('Log-1',10,'-') # -----Log-1 .</code>	LPAD and RPAD — Returns expr1, left or right padded to length n characters with the sequence of characters in expr2: <code>LPAD('Log-1',10,'-') # -----Log-1 .</code>	Yes
REGEXP_REPLACE — Search a string for a regular expression pattern: <code>regexp_replace('John', '[hn].', '1') # Jo1.</code>	You can simulate Oracle REGEXP_REPLACE function using MySQL built-in function.	No
REGEXP_SUBSTR — Extends the functionality	You can simulate Oracle REGEXP_SUBSTR function	No

Oracle function and definition	MySQL function and definition	Equivalent
<p>of the SUBSTR function by searching a string for a regular expression pattern:</p> <pre>REGEXP_SUBSTR('http://www.aws.com/products', 'http://([[a1num:]]+\.\.?)\{3,4\}/?')</pre> <p># http://www.aws.com / .</p>	<p>using MySQL built-in function.</p>	
<p>REPLACE — Returns char with every occurrence of search string replaced with a replacement string: <code>replace('abcdef', 'abc', '123')</code> # 123def.</p>	<p>REPLACE — Returns char with every occurrence of search string replaced with a replacement string: <code>replace('abcdef', 'abc', '123')</code> # 123def.</p>	Yes
<p>LTRIM and RTRIM — Removes from the left or right end of char all of the characters that appear in set: <code>ltrim('zzzyaws', 'xyz')</code> # aws.</p>	<p>LTRIM and RTRIM — Removes spaces from the left or right end of char: <code>ltrim(' Amazon')</code> # Amazon. Combine with the REPLACE function to get the results similar to Oracle.</p>	Partly
<p>SUBSTR — Returns a portion of char, beginning at character position, substring length characters long: <code>substr('John Smith', 6, 1)</code> # S.</p>	<p>SUBSTR — Returns a portion of char, beginning at character position, substring length characters long: <code>substr('John Smith', 6, 1)</code> # S.</p>	Yes

Oracle function and definition	MySQL function and definition	Equivalent
<p>TRIM — Trim leading or trailing characters or both from a character string: <code>trim (both 'x' FROM 'xJohnxx') # John .</code></p>	<p>TRIM — Trim leading or trailing characters or both from a character string: <code>trim (both 'x' FROM 'xJohnxx') # John .</code></p>	Yes
<p>ASCII — Returns the decimal representation in the database character set of the first character of char: <code>ascii('a') # 97 .</code></p>	<p>ASCII — Returns the decimal representation in the database character set of the first character of char: <code>ascii('a') # 97 .</code></p>	Yes
<p>INSTR — Search string for substring.</p>	<p>INSTR — Search string for substring.</p>	Yes
<p>LENGTH — Returns the length of char: <code>length ('John S.') # 7.</code></p>	<p>LENGTH — Returns the length of char: <code>length ('John S.') # 7.</code></p>	Yes
<p>REGEXP_COUNT — Returns the number of times, a pattern occurs in a source string.</p>	<p>You can simulate Oracle REGEXP_COUNT function using MySQL built-in function.</p>	No
<p>REGEXP_INSTR — Searches a string position for a regular expression pattern.</p>	<p>You can simulate Oracle REGEXP_INSTR function using MySQL built-in function.</p>	No

Date and time functions

Oracle function and definition	MySQL function and definition	Equivalent
<p>ADD_MONTHS — Returns the date plus integer months:</p> <pre>add_months(sysdate, 1)</pre>	<p>ADDDATE — MySQL can implement the same functionality using the ADDDATE function.</p>	No
<p>CURRENT_DATE — Returns the current date in the session time zone:</p> <pre>select current_date from dual # 2017-01-01 13:01:01.</pre>	<p>CURRENT_DATE — Returns date without time. Use the now() or the current_timestamp function to achieve the same results:</p> <pre>select now() # 2017-01-01 13:01:01 .</pre>	Partly
<p>CURRENT_TIMESTAMP — Returns the current date and time in the session time zone:</p> <pre>select current_timestamp from dual; # 2017-01-01 13:01:01 .</pre>	<p>CURRENT_TIMESTAMP — Returns the current date and time in the session time zone:</p> <pre>select current_timestamp from dual; # 2017-01-01 13:01:01 .</pre>	Yes
<p>EXTRACT (date part) — Returns the value of a specified date time field from a date time or interval expression:</p> <pre>EXTRACT (YEAR FROM DATE '2017-03-07') # 2017 .</pre>	<p>EXTRACT (date part) — Returns the value of a specified date time field from a date time or interval expression:</p> <pre>EXTRACT (YEAR FROM DATE '2017-03-07') # 2017 .</pre>	Yes
<p>LAST_DAY — Returns the date of the last day of the month that contains date:</p>	<p>LAST_DAY — Returns the date of the last day of the month that contains date:</p>	Yes

Oracle function and definition	MySQL function and definition	Equivalent
<p>LAST_DAY('05-07-2018') # 05-31-2018 .</p>	<p>LAST_DAY('05-07-2018') # 05-31-2018 .</p>	
<p>BETWEEN — Returns the number of months between dates date1 and date2: MONTHS_BETWEEN (sysdate, sysdate-100) # 3.25 .</p>	<p>PERIOD_DIFF — Returns the number of months between periods P1 and P2. P1 and P2 should be in the format YYMM or YYYYMM: SELECT PERIOD_DIFF(201801,201703) # 10</p>	Partly
<p>SYSDATE — Returns the current date and time set for the operating system on which the database server resides: select sysdate from dual # 2017-01-01 13:01:01 .</p>	<p>SYSDATE — Returns the current date and time set for the operating system on which the database server resides: select sysdate() # 2017-01-01 13:01:01 .</p>	Yes
<p>SYSTIMESTAMP — Returns the system date, including fractional seconds and time zone: select systimestamp from dual # 2017-01-01 13:01:01.123456 PM+00:00 .</p>	<p>CURRENT_TIMESTAMP — Returns the current date and time in the session time zone: select current_timestamp from dual; # 2017-01-01 13:01:01.123456+00 .</p>	Yes

Oracle function and definition	MySQL function and definition	Equivalent
<p>LOCALTIMESTAMP — Returns the current date and time in the session time zone in a value of the TIMESTAMP data type:</p> <pre>select localtimestamp from dual # 01-JAN-17 10.01.10.123456 PM .</pre>	<p>LOCALTIMESTAMP — Returns the current date and time in the session time zone in a value of the TIMESTAMP data type:</p> <pre>select localtimestamp from dual # 01-JAN-17 10.01.10.123456 PM .</pre>	Yes
<p>TO_CHAR(datetime) — Converts a date time or timestamp to data type to a value of VARCHAR2 data type in the format specified by the date format: to_char(sys-date, 'DD-MON-YYYY HH24:MI:SS') # 01-JAN-2017 01:01:01 .</p>	<p>DATE_FORMAT — Changes the format of the date and time: DATE_FORMAT(SYSDATE(), '%Y-%m-%d %H:%i:%s')</p>	Yes
<p>TRUNC (date) — Returns a date with the time portion of the day truncated to the unit specified by the format model: trunc(systimestamp) # 2017-01-01 00:00:00 .</p>	<p>You can simulate Oracle TRUNC function using MySQL built-in function.</p>	No

Encoding and decoding functions

Oracle function and definition	MySQL function and definition	Equivalent
DECODE — Compares an expression to each search value one by one using the functionality of an IF-THEN-ELSE statement.	CASE — Compares an expression to each search value one by one.	No
DUMP — Returns a VARCHAR2 value containing the data type code, length in bytes, and internal representation of expression.	N/A	No
ORA_HASH — Computes a hash value for a given expression.	SHA — Calculates an SHA-1 160-bit checksum for the string.	No

Null functions

Oracle function and definition	MySQL function and definition	Equivalent
CASE — Chooses from a sequence of conditions and runs a corresponding statement: CASE WHEN condition THEN result [WHEN ...] [ELSE result] END.	CASE — Chooses from a sequence of conditions and runs a corresponding statement: CASE WHEN condition THEN result [WHEN ...] [ELSE result] END.	Yes
COALESCE — Returns the first non-null expr in the expression	COALESCE — Returns the first of its arguments that	Yes

Oracle function and definition	MySQL function and definition	Equivalent
<p><code>coalesce (null, 'a', 'b') # a.</code></p> <p>NULLIF — Compares <code>expr1</code> and <code>expr2</code>. If they are equal, the function returns null. If they aren't equal, the function returns <code>expr1</code>: <code>NULLIF('a', 'b') # a .</code></p>	<p>isn't null: <code>coalesce (null, 'a', 'b') # a.</code></p> <p>NULLIF — Compares <code>expr1</code> and <code>expr2</code>. If they are equal, the function returns null. If they aren't equal, the function returns <code>expr1</code>: <code>NULLIF('a', 'b') # a .</code></p>	Yes
<p>NVL — Replaces null (returned as a blank) with a string in the results of a query: <code>NVL (null, 'a') # a.</code></p>	<p>IFNULL — Replaces null (returned as a blank) with a string in the results of a query: <code>IFNULL (null, 'a') # a.</code></p>	No
<p>NVL2 — Determines the value returned by a query based on whether a specified expression is null or not null.</p>	<p>CASE — Chooses from a sequence of conditions and runs a corresponding statement: <code>CASE WHEN condition THEN result [WHEN ...] [ELSE result] END.</code></p>	No

Environment and identifier functions

Oracle function and definition	MySQL function and definition	Equivalent
<p><code>SYS_GUID</code> — Generates and returns a globally unique identifier (RAW value) made up of 16 bytes: <code>select sys_guid() from dual</code></p>	<p><code>UUID</code> and <code>REPLACE</code> — <code>REPLACE(UUID(), '-', '')</code>.</p>	No

Oracle function and definition	MySQL function and definition	Equivalent
# 5A280ABA8C76201EE0 530-100007FF691 .		
UID — Returns an integer that uniquely identifies the session user (the user who logged on): <code>select uid from dual # 84.</code>	N/A	No
USER — Returns the name of the session user: <code>select user from dual.</code>	USER — Returns the name of the session user and source machine: <code>select USER().</code>	No
USERENV — Returns information about the current session using parameters: <code>SELECT USERENV ('LANGUAGE') "Language" FROM DUAL .</code>	SHOW SESSION VARIABLES — Displays the system variable values that are in effect for the current connection: <code>myshow SESSION VARIABLES LIKE 'collation_connection'; .</code>	No

Oracle conversion functions

Oracle function and definition	MySQL function and definition	Equivalent
CAST — Converts one built-in data type or collection-typed value into another built-in data type or collection-typed value: <code>cast ('10' as int) + 1 # 11.</code>	CAST — Converts one built-in data type or collection-typed value into another built-in data type or collection-typed value: <code>cast ('10' as UNSIGNED) + 1.</code>	Yes

Oracle function and definition	MySQL function and definition	Equivalent
<p>CONVERT — Converts a character string from a one-character set to another:</p> <pre>select convert ('Ä Ê Í Õ Ø A B C D E ', 'US7ASCII', 'WE8ISO8859P1') from dual .</pre>	<p>CONVERT — Converts a character string from a one-character set to another:</p> <pre>select convert ('Ä Ê Í Õ Ø A B C D E ' USING utf8).</pre>	Yes
<p>TO_CHAR (string / numeric) — Converts NCHAR, NVARCHAR2 , CLOB, or NCLOB data to the database character set:</p> <pre>select to_char ('01234') from dual # 01234.</pre>	<p>FORMAT — Converts string data to the database character set:</p> <pre>FORMAT('01234', 0) -# 01234 .</pre>	No
<p>TO_DATE — Converts char of CHAR, VARCHAR2, NCHAR, or NVARCHAR2 data type to a value of DATE data type:</p> <pre>to_date('01Jan2017', 'DDMonYYYY') # 01-JAN-17 .</pre>	<p>STR_TO_DATE — Convert string data type to a value of DATE data type:</p> <pre>SELECT STR_TO_DATE('01Jan2017', '%d%M%Y') .</pre>	No
<p>TO_NUMBER — Converts an expression to a value of NUMBER data type:</p> <pre>to_number('01234') # 1234 or to_number('01234', '99999') # 1234.</pre>	N/A	No

Aggregate functions

Oracle function and definition	MySQL function and definition	Equivalent
AVG — Returns an average value of an expression: <code>select avg(salary) from employees .</code>	AVG — Returns an average value of an expression: <code>select avg(salary) from employees .</code>	Yes
COUNT — Returns the number of rows returned by the query: <code>select count(*) from employees .</code>	COUNT — Returns the number of rows returned by the query: <code>select count(*) from employees .</code>	Yes
LISTAGG — Orders data within each group specified in the ORDER BY clause and then concatenates the values of the measure column: <code>select listagg(firstname, ' , ') within group (order by customerid) from customer.</code>	GROUP_CONCAT — Orders data within each group specified in the ORDER BY clause and then concatenates the values of the measure column: <code>select GROUP_CONCAT(firstname order by customerid) from customer.</code>	No
MAX — Returns the maximum value of an expression: <code>select max(salary) from employees .</code>	MAX — Returns the maximum value of an expression: <code>select max(salary) from employees .</code>	Yes
MIN — Returns the minimum value of an expression: <code>select min(salary) from employees .</code>	MIN — Returns the minimum value of an expression: <code>select min(salary) from employees .</code>	Yes
SUM — Returns the sum of values of an expression:	SUM — Returns the sum of values of an expression:	Yes

Oracle function and definition	MySQL function and definition	Equivalent
<code>select sum(salary) from employees .</code>	<code>select sum(salary) from employees .</code>	

Top-N Query Oracle 12c

Oracle function and definition	MySQL function and definition	Equivalent
FETCH — Retrieves rows of data from the result set of a multi-row query: <code>select * from customer fetch first 10 rows only.</code>	LIMIT — Retrieves just a portion of the rows that are generated by the rest of the query: <code>select * from customer LIMIT 10.</code>	Yes

For more information, see [String Functions and Operators](#) and [Numeric Functions and Operators](#) in the *MySQL documentation*.

Oracle and MySQL CREATE TABLE AS SELECT statement

With Amazon DMS, you can create a new table in a target database by selecting data from one or more tables in a source database using the Oracle and MySQL CREATE TABLE AS SELECT statement. This statement defines a new table by querying data from existing tables, providing a way to replicate table structures and data from a source to a target database.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	N/A

Oracle usage

The Create Table As Select (CTAS) statement creates a new table based on an existing table. It copies the table DDL definitions (column names and column datatypes) and data to a new table. The new table is populated from the columns specified in the SELECT statement, or all columns if you use SELECT * FROM. You can filter specific data using the WHERE and AND statements. Additionally, you can create a new table having a different structure using joins, GROUP BY, and ORDER BY.

Examples

The following example creates a table based on an existing table and include data from all columns.

```
CREATE TABLE EMPS
AS
SELECT * FROM EMPLOYEES;
```

The following example creates a table based on an existing table with select columns.

```
CREATE TABLE EMPS
AS
SELECT EMPLOYEE_ID, FIRST_NAME, SALARY FROM EMPLOYEES
ORDER BY 3 DESC
```

For more information, see [CREATE TABLE](#) in the *Oracle documentation*.

MySQL usage

MySQL conforms to the ANSI/SQL standard for CTAS functionality and is compatible with an Oracle CTAS statement. For MySQL, the following CTAS standard elements are optional:

- The standard requires parentheses around the SELECT statement; MySQL doesn't.
- The standard requires the WITH [NO] DATA clause; MySQL doesn't.

Examples

The following example creates a table based on an existing table and include data from all columns.

```
CREATE TABLE EMPS AS SELECT * FROM EMPLOYEES;
```

The following example creates a table based on an existing table with select columns.

```
CREATE TABLE EMPS AS SELECT EMPLOYEE_ID, FIRST_NAME, SALARY FROM EMPLOYEES ORDER BY 3
DESC;
```

Oracle and MySQL Common Table Expressions

The following sections provide details on defining and leveraging Common Table Expressions (CTEs) within Amazon DMS to streamline database operations and enhance query performance.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Common Table Expressions	MySQL doesn't support common table expressions. A workaround is available.

Oracle usage

CTEs provide a way to implement the logic of sequential code or to reuse code. You can define a named sub query and then use it multiple times in different parts of a query statement.

A CTE is implemented using a WITH clause, which is part of the ANSI SQL-99 standard and has existed in Oracle since version 9.2. CTE usage is similar to an inline view or a temporary table. Its main purpose is to reduce query statement repetition and make complex queries simpler to read and understand.

Syntax

```
WITH <subquery name> AS (<subquery code>)[...]
```

```
SELECT <Select list> FROM <subquery name>;
```

Examples

The following example creates a sub query of the employee count for each department and then use the result set of the CTE in a query.

```
WITH DEPT_COUNT
(DEPARTMENT_ID, DEPT_COUNT) AS
(SELECT DEPARTMENT_ID, COUNT(*)
FROM EMPLOYEES
GROUP BY DEPARTMENT_ID)
SELECT E.FIRST_NAME || ' ' || E.LAST_NAME AS EMP_NAME,
D.DEPT_COUNT AS EMP_DEPT_COUNT
FROM EMPLOYEES E JOIN DEPT_COUNT D
USING (DEPARTMENT_ID)
ORDER BY 2;
```

MySQL usage

Aurora MySQL 5.7 doesn't support common table expressions (CTE).

Note

Amazon Relational Database Service (Amazon RDS) for MySQL version 8 supports common table expressions both non-recursive and recursive. Common table expressions enable use of named temporary result sets implemented by permitting a WITH clause preceding SELECT statements and certain other statements. For more information, see [WITH \(Common Table Expressions\)](#). As of MySQL 8.0.19, the recursive SELECT part of a recursive common table expression (CTE) supports a LIMIT clause. LIMIT with OFFSET is also supported. For more information, see [Recursive Common Table Expressions](#).

Migration considerations

As a workaround, use views or derived tables in place of non-recursive CTEs. Since non-recursive CTEs are more convenient for readability and code simplification, you can convert the code to use derived tables, which are a subquery in the parent query's FROM clause. For example, replace the following CTE:

```
WITH TopCustomerOrders
( SELECT Customer, COUNT(*) AS NumOrders
  FROM Orders
  GROUP BY Customer
)
SELECT TOP 10 * FROM TopCustomerOrders ORDER BY NumOrders DESC;
```

With the following subquery:

```
SELECT *
FROM ( SELECT Customer, COUNT(*) AS NumOrders
      FROM Orders
      GROUP BY Customer ) AS TopCustomerOrders
ORDER BY NumOrders DESC
LIMIT 10 OFFSET 0;
```

When you use derived tables, make sure that the derived table definition is repeated if multiple instances are required for the query.

Converting the code for recursive CTEs is not straight forward, but you can achieve similar functionality using loops.

Examples

Replacing non-recursive CTEs

Use a derived table to replace non-recursive CTE functionality as follows:

The following example creates and populates an OrderItems table.

```
CREATE TABLE OrderItems(
  OrderID INT NOT NULL, Item VARCHAR(20) NOT NULL,
  Quantity SMALLINT NOT NULL,
  PRIMARY (OrderID, Item));

INSERT INTO OrderItems (OrderID, Item, Quantity)
VALUES (1, 'M8 Bolt', 100), (2, 'M8 Nut', 100),
(3, 'M8 Washer', 200), (3, 'M6 Washer', 100);
```

Define a derived table for TotalQty of every order and then join to the OrderItems to obtain the relative quantity for each item.

```

SELECT O.OrderID, O.Item, O.Quantity, (O.Quantity / A0.TotalQty) * 100 AS
  PercentOfOrder
FROM OrderItems AS O
  INNER JOIN
  ( SELECT OrderID, SUM(Quantity) AS TotalQty
    FROM OrderItems
    GROUP BY OrderID
  ) AS A0
ON O.OrderID = A0.OrderID;

```

For the preceding example, the result looks as shown following.

OrderID	Item	Quantity	PercentOfOrder
1	M8 Bolt	100	100.0000000000
2	M8 Nut	100	100.0000000000
3	M8 Washer	100	33.3333333300
3	M6 Washer	200	66.6666666600

Replacing recursive CTEs

Use recursive SQL code in stored procedures and SQL loops to replace a recursive CTEs.

Note

Stored procedure and function recursion in Aurora MySQL is turned off by default. You can set the server system variable `max_sp_recursion_depth` to a value of 1 or higher to turn on recursion. However, this approach is not recommended because it may increase contention for the thread stack space.

The following example creates and populates an Employees table.

```

CREATE TABLE Employees
( Employee VARCHAR(5) NOT NULL PRIMARY KEY,
  DirectManager VARCHAR(5) NULL);

INSERT INTO Employees (Employee, DirectManager)
VALUES ('John', 'Dave'), ('Jose', 'Dave'),
('Fred', 'John'), ('Dave', NULL);

```

The following example creates an EmpHierarchy table.

```
CREATE TABLE EmpHierarchy (LVL INT, Employee VARCHAR(5), Manager VARCHAR(5));
```

The following example creates a procedure that uses a loop to traverse the employee hierarchy.

```
CREATE PROCEDURE P()  
BEGIN  
  DECLARE var_lvl INT;  
  DECLARE var_Employee VARCHAR(5);  
  SET var_lvl = 0;  
  SET var_Employee = (SELECT Employee FROM Employees WHERE DirectManager IS NULL);  
  INSERT INTO EmpHierarchy VALUES (var_lvl, var_Employee, NULL);  
  WHILE var_lvl <> -1  
  DO  
    INSERT INTO EmpHierarchy (LVL, Employee, Manager)  
    SELECT var_lvl + 1, Employee, DirectManager  
    FROM Employees  
    WHERE DirectManager IN (SELECT Employee FROM EmpHierarchy WHERE LVL = var_lvl);  
    IF NOT EXISTS (SELECT * FROM EmpHierarchy WHERE LVL = var_lvl + 1)  
    THEN SET var_lvl = -1;  
    ELSE SET var_lvl = var_lvl + 1;  
    END IF;  
  END WHILE;  
END;
```

Run the procedure.

```
CALL P();
```

Select all records from the EmpHierarchy table.

```
SELECT * FROM EmpHierarchy;
```

Level	Employee	Manager
0	Dave	
1	John	Dave
1	Jose	Dave
2	Fred	John

Summary

Oracle	Aurora MySQL	Comments
Non-recursive CTE	Derived table	For multiple instances of the same table, the derived table definition subquery must be repeated.
Recursive CTE	Loop inside a stored procedure or stored function.	

For more information, see [WITH \(Common Table Expressions\)](#) in the *MySQL documentation*.

Oracle sequences and identity columns and MySQL sequences and AUTO INCREMENT columns

Oracle sequences and identity columns, as well as MySQL Sequences and AUTO_INCREMENT columns, are database objects used to generate unique sequential values, often employed as primary keys or unique identifiers. The following sections provide detailed guidance on handling Oracle sequences and identity columns, and MySQL sequences and AUTO_INCREMENT columns when using Amazon DMS.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Sequences	MySQL doesn't support sequences, identity columns have different syntax and options.

Oracle usage

Sequences are database objects that serve as unique identity value generators. You can use them, for example, to automatically generate primary key values. Oracle treats sequences as independent objects. The same sequence can generate values for multiple tables.

You can configure sequences with multiple parameters to control their value-generating behavior. For example, the `INCREMENT BY` sequence parameter defines the interval between each generated sequence value. If more than one database user is generating incremented values from the same sequence, each user may encounter gaps in the generated values that are visible to them.

Oracle 18c introduced scalable sequences: a special class of sequences that are optimized for multiple concurrent session usage.

This introduces three new options when creating a new sequence:

- `SCALE` — Turns on the sequence scalability feature.
 - `EXTEND` — Extends in additional 6 digits offset (as default) and the maximum number of digits in the sequence (`maxvalue` and `minvalue`).
 - `NOEXTEND` — sequence value will be padded to the max value. This is the default option when using the `SCALE` option.
- `NOSCALE` - non-scalable sequence usage.

Oracle sequence options

By default, the initial and increment values for a sequence are both 1, with no upper limit.

- `INCREMENT BY` — Controls the sequence interval value of the increment or decrement (if a negative value is specified). If the `INCREMENT BY` parameter isn't specified during sequence creation, the value is set to 1. The increment cannot be assigned a value of 0.
- `START WITH` — Defines the initial value of a sequence. The default value is 1.
- `MAXVALUE` and `NOMAXVALUE` — Specifies the maximum limit for values generated by a sequence. It must be equal or greater than the `START WITH` parameter and must be greater in value than the `MINVALUE` parameter. The default for `NOMAXVALUE` is 1027 for an ascending sequence.
- `MINVALUE` and `NOMINVALUE` — Specifies the minimum limit for values generated by a sequence. Must be less than or equal to the `START WITH` parameter and must be less than the `MAXVALUE` parameter. The default for `NOMINVALUE` is -1026 for a descending sequence.

- **CYCLE and NOCYCLE** — Instructs a sequence to continue generating values despite reaching the maximum value or the minimum value. If the sequence reaches one of the defined ascending limits, it generates a new value according to the minimum value. If it reaches a descending limit, it generates a new value according to the maximum value. The default option is NOCYCLE.
- **CACHE and NOCACHE** — Specifies the number of sequence values to keep cached in memory for improved performance. CACHE has a minimum value of 2. The NOCACHE parameter causes a sequence to not cache values in memory. Specifying neither CACHE nor NOCACHE will cache 20 values to memory. In the event of a database failure, all unused cached sequence values are lost and gaps in sequence values may occur.
- **SCALE and NOSCALE**: Turns on the scalable sequences feature.

Examples

Create a sequence.

```
CREATE SEQUENCE SEQ_EMP
START WITH 100
INCREMENT BY 1
MAXVALUE 9999999999
CACHE 20
NOCYCLE;
```

Drop a sequence.

```
DROP SEQUENCE SEQ_EMP;
```

View sequences created for the current schema or user.

```
SELECT * FROM USER_SEQUENCES;
```

Use a sequence as part of an INSERT INTO statement.

```
CREATE TABLE EMP_SEQ_TST (COL1 NUMBER PRIMARY KEY, COL2 VARCHAR2(30));
INSERT INTO EMP_SEQ_TST VALUES(SEQ_EMP.NEXTVAL, 'A');

COL1  COL2
100   A
```

Query the current value of a sequence.

```
SELECT SEQ_EMP.CURRVAL FROM DUAL;
```

Manually increment the value of a sequence according to the INCREMENT BY specification.

```
SELECT SEQ_EMP.NEXTVAL FROM DUAL;
```

Alter an existing sequence.

```
ALTER SEQUENCE SEQ_EMP MAXVALUE 1000000;
```

Create a scalable sequence.

```
CREATE SEQUENCE scale_seq
MINVALUE 1
MAXVALUE 9999999999
SCALE;

select scale_seq.nextval as scale_seq from dual;

NEXTVAL
1010320001
```

Oracle 12c default values using sequences

Starting from Oracle 12c, you can assign a sequence to a table column with the CREATE TABLE statement and specify the NEXTVAL configuration of the sequence.

Generate DEFAULT values using sequences.

```
CREATE TABLE SEQ_TST ( COL1 NUMBER DEFAULT SEQ_1.NEXTVAL PRIMARY KEY, COL2
  VARCHAR(30));

INSERT INTO SEQ_TST(COL2) VALUES('A');

SELECT * FROM SEQ_TST;

COL1 COL2
100 A
```

Oracle 12c session sequences

Starting from Oracle 12c, you can create sequences as session-level or global-level. By adding the `SESSION` parameter to a `CREATE SEQUENCE` statement, the sequence is created as a session-level sequence. Optionally, the `GLOBAL` keyword can be used to create a global sequence to provide consistent results across sessions in the database. Global sequences are the default. Session sequences return a unique range of sequence numbers only within a session.

The following example creates Oracle 12c `SESSION` and `GLOBAL` sequences.

```
CREATE SEQUENCE SESSION_SEQ SESSION;  
CREATE SEQUENCE SESSION_SEQ GLOBAL;
```

Oracle 12c identity columns

Oracle 12c introduced support for automatic generation of values to populate columns in database tables. The `IDENTITY` type generates a sequence and associates it with a table column without the need to manually create a separate Sequence object. It relies internally on sequences and can be manually configured.

Sequences can be used as an `IDENTITY` type, which automatically creates a sequence and associates it with the table column. The main difference is that there is no need to create a sequence manually; the `IDENTITY` type does that for you. An `IDENTITY` type is a sequence that can be configured.

Create a table with an Oracle 12c Identity Column.

```
CREATE TABLE IDENTITY_TST (  
  COL1 NUMBER GENERATED BY DEFAULT AS IDENTITY (START WITH 100  
  INCREMENT BY 10),  
  COL2 VARCHAR2(30));
```

Insert records using an Oracle 12c `IDENTITY` column (explicitly/implicitly).

```
INSERT INTO IDENTITY_TST(COL2) VALUES('A');  
INSERT INTO IDENTITY_TST(COL1, COL2) VALUES(DEFAULT, 'B');  
INSERT INTO IDENTITY_TST(co11, co12) VALUES(NULL, 'C');  
  
SELECT * FROM IDENTITY_TST;
```

```
COL1  COL2
120   A
130   B
```

For more information, see [CREATE SEQUENCE](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL supports automatic sequence generation using the `AUTO_INCREMENT` column property, similar to the Oracle `IDENTITY` column property. It doesn't support table-independent sequence objects.

Any numeric column may be assigned the `AUTO_INCREMENT` property. To make the system generate the next sequence value, the application must not mention the relevant column's name in the insert command, in case the column was created with the `NOT NULL` definition then also inserting a `NULL` value into an `AUTO_INCREMENT` column will increment it. In most cases, the seed value is 1 and the increment is 1.

Client applications use the `LAST_INSERT_ID` function to obtain the last generated value.

Each table can have only one `AUTO_INCREMENT` column. Make sure that the column is explicitly indexed or is a primary key (which is indexed by default).

The `AUTO_INCREMENT` mechanism is designed to be used with positive numbers only. Do not use negative values because they are misinterpreted as a complementary positive value. This limitation is due to precision issues with sequences crossing a zero boundary.

There are two server parameters used to alter the default values for new `AUTO_INCREMENT` columns:

- `auto_increment_increment` — Controls the sequence interval.
- `auto_increment_offset` — Determines the starting point for the sequence.

To reseed the `AUTO_INCREMENT` value, use `ALTER TABLE <Table Name> AUTO_INCREMENT = <New Seed Value>`.

Migration considerations

Because Aurora MySQL doesn't support table-independent `SEQUENCE` objects, applications that rely on its properties must use custom solutions to meet their requirements.

You can use Aurora MySQL `AUTO_INCREMENT` instead of Oracle `IDENTITY` for most cases. For `AUTO_INCREMENT` columns, the application must explicitly `INSERT` a `NULL` or a `0`.

Note

Omitting the `AUTO_INCREMENT` column from the `INSERT` column list has the same effect as inserting a `NULL` value.

Make sure that the `AUTO_INCREMENT` columns are indexed (the following section explains why) and cannot have default constraints assigned to the same column. There is a critical difference between `IDENTITY` and `AUTO_INCREMENT` in the way the sequence values are maintained upon service restart. Application developers must be aware of this difference.

Sequence value initialization

Oracle stores the `IDENTITY` metadata in system tables on disk. Although some values may be cached and are lost when the service is restarted, the next time the server restarts, the sequence value continues after the last block of values that was assigned to cache. If you run out of values, you can explicitly set the sequence value to start the cycle over. As long as there are no key conflicts, it can be reused after the range has been exhausted.

In Aurora MySQL, an `AUTO_INCREMENT` column for a table uses a special auto-increment counter to assign new values for the column. This counter is stored in cache memory only and is not persisted to disk. After a service restart, and when Aurora MySQL encounters an `INSERT` to a table that contains an `AUTO_INCREMENT` column, it issues an equivalent to the following statement:

```
SELECT MAX(<Auto Increment Column>) FROM <Table Name> FOR UPDATE;
```

Note

The `FOR UPDATE CLAUSE` is required to maintain locks on the column until the read completes.

Aurora MySQL then increments the value retrieved by the statement above and assigns it to the in-memory autoincrement counter for the table.

By default, the value is incremented by one. You can change this default using the `auto_increment_increment` configuration setting. If the table has no values, Aurora MySQL uses the value 1. You can change the default using the `auto_increment_offset` configuration setting.

Every server restart effectively cancels any `AUTO_INCREMENT = <Value>` table option in `CREATE TABLE` and `ALTER TABLE` statements.

Unlike Oracle `IDENTITY` columns, which by default do not allow inserting explicit values, Aurora MySQL allows explicit values to be set. If a row has an explicitly specified `AUTO_INCREMENT` column value and the value is greater than the current counter value, the counter is set to the specified column value.

Examples

Create a table with an `AUTO_INCREMENT` column.

```
CREATE TABLE MyTable (Col1 INT NOT NULL AUTO_INCREMENT PRIMARY KEY,  
Col2 VARCHAR(20) NOT NULL);
```

Insert `AUTO_INCREMENT` values.

```
INSERT INTO MyTable (Col2) VALUES ('AI column omitted');
```

```
INSERT INTO MyTable (Col1, Col2) VALUES (NULL, 'Explicit NULL');
```

```
INSERT INTO MyTable (Col1, Col2) VALUES (10, 'Explicit value');
```

```
INSERT INTO MyTable (Col2) VALUES ('Post explicit value');
```

```
SELECT * FROM MyTable;
```

```
Col1  Col2  
1     AI column omitted  
2     Explicit NULL  
10    Explicit value  
11    Post explicit value
```

Reseed AUTO_INCREMENT.

```
ALTER TABLE MyTable AUTO_INCREMENT = 30;
```

```
INSERT INTO MyTable (Col2) VALUES ('Post ALTER TABLE');
```

```
SELECT * FROM MyTable;
```

Col1	Col2
1	AI column omitted
2	Explicit NULL
10	Explicit value
11	Post explicit value
30	Post ALTER TABLE

Summary

The following table identifies similarities, differences, and key migration considerations.

Feature	Oracle	Aurora MySQL	Comments
Create a table	<pre>CREATE TABLE IDENTITY_TST (COL1 NUMBER GENERATED BY DEFAULT AS IDENTITY (START WITH 100 INCREMENT BY 10), COL2 VARCHAR2(30));</pre>	<pre>CREATE TABLE AUTO_TST (COL1 INT AUTO_INCREMENT PRIMARY KEY, COL2 VARCHAR(3 0));</pre>	
Set the starting number	<pre>CREATE TABLE IDENTITY_TST (COL1 NUMBER GENERATED BY DEFAULT</pre>	<pre>ALTER TABLE AUTO_TST AUTO_INCREMENT = 100;</pre>	

Feature	Oracle	Aurora MySQL	Comments
	<pre>AS IDENTITY (START WITH 100 INCREMENT BY 10), COL2 VARCHAR2(30));</pre>	Or use the <code>auto_increment_offset</code> parameter	
Set the interval	<pre>CREATE TABLE IDENTITY_TST (COL1 NUMBER GENERATED BY DEFAULT AS IDENTITY (START WITH 100 INCREMENT BY 10), COL2 VARCHAR2(30));</pre>	Set the <code>auto_increment_increment</code> parameter	
Additional permitted values	DEFAULT, NULL	None	
Independent SEQUENCE object	CREATE SEQUENCE	Not supported	
Automatic enumerator column property	IDENTITY	AUTO_INCREMENT	
Reseed sequence value	Recreate the sequence	ALTER TABLE	
Column restrictions	Numeric	Numeric, indexed, and no DEFAULT	

Feature	Oracle	Aurora MySQL	Comments
Controlling seed and interval values	CREATE/ALTER TABLE	auto_increment and auto_increment_offset	Aurora MySQL settings are global and can't be customized for each column as with Oracle.
Sequence setting initialization	Maintained through service restarts	Re-initialized every service restart	For more information, see Sequence Value Initialization .
Explicit values to column	Not supported	Supported	Aurora MySQL requires explicit NULL or 0 to trigger sequence value assignment. Inserting an explicit value larger than all others reinitializes the sequence.

For more information, see [Using AUTO_INCREMENT](#), [CREATE TABLE Statement](#), and [InnoDB AUTO_INCREMENT Counter Initialization](#) in the *MySQL documentation*.

Oracle and MySQL INSERT FROM SELECT statement

The following sections provide details on running the INSERT FROM SELECT statement, including syntax examples and best practices for efficient data transfer.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	MySQL doesn't support ERROR LOG and subquery options.

Oracle usage

You can insert multiple records into a table from another table using the `INSERT FROM SELECT` statement, which is a derivative of the basic `INSERT` statement. The column ordering and data types must match between the target and the source tables.

Examples

Simple `INSERT FROM SELECT` (explicit).

```
INSERT INTO EMPS (EMPLOYEE_ID, FIRST_NAME, SALARY, DEPARTMENT_ID) SELECT EMPLOYEE_ID,
FIRST_NAME, SALARY, DEPARTMENT_ID
FROM EMPLOYEES
WHERE SALARY > 10000;
```

Simple `INSERT FROM SELECT` (implicit).

```
INSERT INTO EMPS
SELECT EMPLOYEE_ID, FIRST_NAME, SALARY, DEPARTMENT_ID
FROM EMPLOYEES
WHERE SALARY > 10000;
```

This example produces the same result as the preceding example but uses a subquery in the `DML_table_expression_clause`.

```
INSERT INTO
(SELECT EMPLOYEE_ID, FIRST_NAME, SALARY, DEPARTMENT_ID FROM EMPS)
VALUES (120, 'Kenny', 10000, 90);
```

Log errors with the Oracle `error_logging_clause`.

```
ALTER TABLE EMPS ADD CONSTRAINT PK_EMP_ID PRIMARY KEY(employee_id);
EXECUTE DBMS_ERRLOG.CREATE_ERROR_LOG('EMPS', 'ERRLOG');
INSERT INTO EMPS
SELECT EMPLOYEE_ID, FIRST_NAME, SALARY, DEPARTMENT_ID
FROM EMPLOYEES
WHERE SALARY > 10000
LOG ERRORS INTO errlog ('Cannot Perform Insert') REJECT LIMIT 100;
0 rows inserted
```

When inserting an existing `EMPLOYEE ID` into the `EMPS` table, the insert doesn't fail because the invalid records are redirected to the `ERRLOG` table.

For more information, see [INSERT](#) in the *Oracle documentation*.

MySQL usage

MySQL is compatible with the Oracle `INSERT FROM SELECT` syntax except for a few features specific to Oracle. For example, the `conditional_insert_clause` (`ALL | FIRST | ELSE`). MySQL doesn't support the Oracle `error_logging_clause` feature. Generally, you can use `ON DUPLICATE KEY UPDATE` to handle duplicate rows.

Syntax

```
INSERT [LOW_PRIORITY | HIGH_PRIORITY] [IGNORE]
  [INTO] tbl_name
  [PARTITION (partition_name [, partition_name] ...)]
  [(col_name [, col_name] ...)]
  SELECT ...
  [ON DUPLICATE KEY UPDATE assignment_list]
```

value:

```
{expr | DEFAULT}
```

assignment:

```
col_name = value
```

assignment_list:

```
assignment [, assignment] ...
```

Examples

Simple INSERT FROM SELECT (explicit).

```
INSERT INTO EMPS (EMPLOYEE_ID, FIRST_NAME, SALARY, DEPARTMENT_ID)
SELECT EMPLOYEE_ID, FIRST_NAME, SALARY, DEPARTMENT_ID
FROM EMPLOYEES
WHERE SALARY > 10000;
```

Simple INSERT FROM SELECT (implicit).

```
INSERT INTO EMPS
SELECT EMPLOYEE_ID, FIRST_NAME, SALARY, DEPARTMENT_ID
FROM EMPLOYEES
WHERE SALARY > 10000;
```

The following example isn't compatible with MySQL.

```
INSERT INTO
(SELECT EMPLOYEE_ID, FIRST_NAME, SALARY, DEPARTMENT_ID FROM EMPS)
VALUES (120, 'Kenny', 10000, 90);
```

The following example demonstrates using the ON DUPLICATE KEY UPDATE clause to update specific columns when a UNIQUE violation occurs.

```
INSERT INTO EMPS
SELECT * from EMPLOYEES
  where EMPLOYEE_ID > 10
ON DUPLICATE KEY UPDATE
EMPS.FIRST_NAME=EMPLOYEES.FIRST_NAME,
EMPS.SALARY=EMPLOYEES.SALARY;
```

For more information, see [INSERT ... SELECT Statement](#) in the *MySQL documentation*.

Multi-Version Concurrency Control

With Amazon DMS, you can implement Multi-Version Concurrency Control (MVCC) to manage concurrent access to data during database migrations. MVCC is a concurrency control method that

maintains multiple versions of database objects, allowing readers and writers to access the data simultaneously without blocking or causing conflicts.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	N/A

Oracle usage

Two primary lock types exist in Oracle: exclusive locks and share locks, which implement the following high-level locking semantics:

- Writers never block readers.
- Readers never block writers.
- Oracle never escalates locks from row to page and table level, which reduces potential deadlocks.
- In Oracle, users can issue explicit locks on specific tables using the `LOCK TABLE` statement.

Lock types can be divided into four categories: DML locks, DDL locks, Explicit (Manual) data locking, and System locks. The following sections describe each category.

DML Locks

DML locks preserve the integrity of data accessed concurrently by multiple users. DML statements acquire locks automatically both on row and table levels.

- **Row locks or TX** — Obtained on a single row of a table by one the following statements: `INSERT`, `UPDATE`, `DELETE`, `MERGE`, and `SELECT ... FOR UPDATE` . If a transaction obtains a row lock, a table lock is also acquired to prevent DDL modifications to the table that might cause conflicts. The lock exists until the transaction ends with a `COMMIT` or `ROLLBACK`.
- **Table locks or TM** — When performing one of the following DML operations: `INSERT`, `UPDATE`, `DELETE`, `MERGE`, and `SELECT ... FOR UPDATE` , a transaction automatically acquires a table lock

to prevent DDL modifications to the table that might cause conflicts if the transaction did not issue a COMMIT or ROLLBACK.

All table lock types:

- **Row share lock or RS** — Occurs when the transaction holding the lock on the table has locked some rows in the table before updating them.
- **Row Exclusive lock or RX** — Occurs when the transaction holding the lock has updated table rows or used the `SELECT ... FOR UPDATE` command.
- **Share table lock or S** — One transaction locks the table and allows other transactions to query the table (exclude `SELECT ... FOR UPDATE`), it also allows updates only if a single transaction holds the share table lock. Multiple transactions may hold a share table lock concurrently.
- **Share row exclusive table lock or SRX** — Similar to S lock but with this lock, only a single transaction at a time can acquire this lock on a given table.
- **Exclusive table lock or X** — Most restrictive lock type, it allows the transaction that holds the lock an exclusive write access to the table. Only one transaction can obtain an X lock for a table.

The following table provides additional information regarding row and table locks.

Statement	Row locks	Table lock mode	RS	RX	S	SRX	X
<code>SELECT ... FROM table ...</code>	—	none	Y	Y	Y	Y	Y
<code>INSERT INTO table ...</code>	Yes	SX	Y	Y	N	N	N
<code>UPDATE table ...</code>	Yes	SX	Y	Y	N	N	N

Statement	Row locks	Table lock mode	RS	RX	S	SRX	X
MERGE INTO table ...	Yes	SX	Y	Y	N	N	N
DELETE FROM table ...	Yes	SX	Y	Y	N	N	N
SELECT ... FROM table FOR UPDATE OF...	Yes	SX	Y	Y	N	N	N
LOCK TABLE table IN...	—						
ROW SHARE MODE		SS	Y	Y	Y	Y	N
ROW EXCLUSIVI MODE		SX	Y	Y	N	N	N
SHARE MODE		S	Y	N	Y	N	N

Statement	Row locks	Table lock mode	RS	RX	S	SRX	X
SHARE ROW EXCLUSIVI MODE		SSX	Y	N	N	N	N
EXCLUSIVI MODE		X	N	N	N	N	N

DDL Locks

The main purpose of a DDL lock is to protect the definition of a schema object while it is modified by an ongoing DDL operation such as `ALTER TABLE EMPLOYEES ADD <COLUMN>`.

Explicit or manual data locking

Users have the ability to explicitly create locks to achieve transaction-level read consistency for when an application requires transactional exclusive access to a resource without waiting for other transactions to complete. Explicit data locking can be performed at the transaction level or the session level:

- Transaction level
 - `SET TRANSACTION ISOLATION LEVEL`
 - `LOCK TABLE`
 - `SELECT ... FOR UPDATE`
- Session level
 - `ALTER SESSION SET ISOLATION LEVEL`

System locks

System locks include latches, mutexes, and internal locks.

Examples

Explicitly lock data using the LOCK TABLE command.

```
-- Session 1
LOCK TABLE EMPLOYEES IN EXCLUSIVE MODE;
-- Session 2
UPDATE EMPLOYEES
SET SALARY=SALARY+1000
WHERE EMPLOYEE_ID=114;
-- Session 2 waits for session 1 to COMMIT or ROLLBACK
```

Explicitly lock data using the SELECT... FOR UPDATE command. Oracle obtains exclusive row-level locks on all the rows identified by the SELECT FOR UPDATE statement.

```
-- Session 1
SELECT * FROM EMPLOYEES WHERE EMPLOYEE_ID=114 FOR UPDATE;
-- Session 2
UPDATE EMPLOYEES
SET SALARY=SALARY+1000
WHERE EMPLOYEE_ID=114;
-- Session 2 waits for session 1 to COMMIT or ROLLBACK
```

For more information, see [Automatic Locks in DDL Operations](#), [Automatic Locks in DML Operations](#), and [Automatic and Manual Locking Mechanisms During SQL Operations](#) in the *Oracle documentation*.

MySQL usage

When using InnoDB, MySQL provides various lock modes to control concurrent access to data in tables. Data consistency is maintained using a Multi-Version Concurrency Control (MVCC) mechanism. Most MySQL commands automatically acquire locks of appropriate modes to ensure that referenced tables are not dropped or modified in incompatible ways while the command runs.

The MVCC mechanism prevents viewing inconsistent data produced by concurrent transactions performing updates on the same rows. MVCC provides strong transaction isolation for each database session and minimizes lock-contention in multi-user environments.

- Similar to Oracle, MVCC locks acquired for querying (reading) data do not conflict with locks acquired for writing data. Reads never block writes and writes never block reads.

- Similar to Oracle, MySQL does not escalate locks to table-level such as when an entire table is locked for writes when a certain threshold of row locks is exceeded.

InnoDB uses three additional fields for each row:

- `DB_TRX_ID` — Indicates the transaction identifier for the last transaction that inserted or updated the row.
- `DB_ROLL_PTR` — Points to an undo log record written to the rollback segment.
- `DB_ROW_ID` — Contains a row ID that increases monotonically as new rows are inserted.

Implicit and explicit transactions (Auto-commit behavior)

Unlike Oracle, MySQL uses auto-commit for transactions by default. However, there are two options to support explicit transactions, which are similar to the default behavior in Oracle (non-auto-commit).

- Use the `START TRANSACTION` (or `BEGIN TRANSACTION`) statements and then `COMMIT` or `ROLLBACK`.
- Set `AUTOCOMMIT` to `OFF` at the session level.

With explicit transactions:

- Users can explicitly issue a lock similar to the `LOCK TABLE` statement in Oracle.
- `SELECT... FOR UPDATE` is supported.

Unlike Oracle there are only two types of table-level locks when using the `LOCK TABLE` command: read lock and write lock.

Read lock or shared S lock

- The session that holds the lock can only read the table.
- Multiple sessions can acquire a `READ` lock for the table at the same time.
- Other sessions can read the table without explicitly acquiring a `READ` lock.
- For InnoDB tables, `READ LOCAL` is the same as `READ`.

Write lock or exclusive X lock

- The session that holds the lock can read and write the table.
- Only the session that holds the lock can access the table. No other session can access it until the lock is released.
- Lock requests for the table by other sessions block while the WRITE lock is held.
- The LOW_PRIORITY modifier is deprecated and has no effect.

For row-level locking:

- **Intention shared IS lock** — Indicates that a transaction intends to set a shared lock.
- **Intention exclusive IX lock** — Indicates that a transaction intends to set an exclusive lock.

	X	IX	S	IS
X	Not permitted	Not permitted	Not permitted	Not permitted
IX	Not permitted	Permitted	Not permitted	Permitted
S	Not permitted	Not permitted	Permitted	Permitted
IS	Not permitted	Permitted	Permitted	Permitted

Records lock

A record lock is a lock on an index record. For example, the `SELECT id FROM emps WHERE id = 50 FOR UPDATE` query prevents any other transaction from inserting, updating, or deleting rows where the value of `emps.id` is 50.

Record locks always lock index records, even if a table is defined with no indexes. For such cases, InnoDB creates a hidden clustered index and uses it for record locking.

Gaps lock

A gap lock is a lock on a gap between index records, before the first index record, or after the last index record. For example, `SELECT id FROM emps WHERE id BETWEEN 50 and 80`

FOR UPDATE prevents other transactions from inserting a value of 60 into the `emps.id` column whether or not there was already any value in the column because the gaps between all existing values in the range are locked.

Transaction-level locking

- SET TRANSACTION ISOLATION LEVEL
- LOCK TABLE
- SELECT ... FOR UPDATE

Syntax

```
LOCK TABLES
tbl_name [[AS] alias] lock_type [, tbl_name [[AS] alias] lock_type] ...

lock_type:
READ [LOCAL] | [LOW_PRIORITY] WRITE
```

MySQL deadlocks

Deadlocks occur when two or more transactions acquired locks on each other's process resources such as table or row. MySQL can detect deadlocks automatically and resolve the event by aborting one of the transactions and allowing the other transaction to complete.

Examples

Obtain an explicit lock on a table using the `LOCK TABLE` command.

```
-- Session 1
START TRANSACTION;
LOCK TABLE EMPLOYEES IN EXCLUSIVE MODE;

-- Session 2
UPDATE EMPLOYEES
SET SALARY=SALARY+1000
WHERE EMPLOYEE_ID=114;

-- Session 2 waits for session 1 to COMMIT or ROLLBACK
```

Explicit lock by the SELECT... FOR UPDATE command. MySQL obtains exclusive row-level locks on rows referenced by the SELECT FOR UPDATE statement. Make sure that this statement runs inside a transaction.

```
-- Session 1
START TRANSACTION;
SELECT * FROM EMPLOYEES WHERE EMPLOYEE_ID=114 FOR UPDATE;

-- Session 2
UPDATE EMPLOYEES
SET SALARY=SALARY+1000
WHERE EMPLOYEE_ID=114;

-- Session 2 waits for session 1 to COMMIT or ROLLBACK
```

Summary

Description	Oracle	MySQL
Dictionary tables to obtain information about locks	<code>v\$lock;</code> <code>v\$locked_object;</code> <code>v\$session_blockers;</code>	<code>SHOW OPEN TABLES WHERE</code> <code>in_use = 1;</code>
Lock a table	<code>BEGIN;</code> <code>LOCK TABLE employees IN</code> <code>SHARE ROW EXCLUSIVE</code> <code>MODE;</code>	<code>LOCK TABLE employees</code> <code>READ</code>
Explicit locking	<code>SELECT * FROM employees</code> <code>WHERE employee_id=102</code> <code>FOR UPDATE;</code>	<code>SELECT * FROM employees</code> <code>WHERE employee_id=102</code> <code>FOR UPDATE;</code>
Explicit locking, options	<code>SELECT ... FOR UPDATE</code>	<code>SELECT ... FOR UPDATE</code>

For more information, see [InnoDB Multi-Versioning](#), [LOCK TABLES and UNLOCK TABLES Statements](#), and [SET TRANSACTION Statement](#) in the *MySQL documentation*.

Oracle MERGE statement and MySQL equivalent

With Amazon DMS, you can perform Oracle MERGE statements and the MySQL equivalent to conditionally insert, update, or delete rows in a target table based on the results of a join with a source table.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Merge	Aurora MySQL doesn't support the MERGE statement . A workaround is available.

Oracle usage

The MERGE statement provides a means to specify single SQL statements that conditionally perform INSERT, UPDATE, or DELETE operations on a target table—a task that would otherwise require multiple logical statements.

The MERGE statement selects record(s) from the source table and then, by specifying a logical structure, automatically performs multiple DML operations on the target table. Its main advantage is to help avoid the use of multiple inserts, updates or deletes. It is important to note that MERGE is a deterministic statement. That is, once a row has been processed by the MERGE statement, it can't be processed again using the same MERGE statement. MERGE is also sometimes known as UPSERT.

Examples

Use MERGE to insert or update employees who are entitled to a bonus (by year).

```
CREATE TABLE EMP_BONUS(EMPLOYEE_ID NUMERIC, BONUS_YEAR VARCHAR2(4),
SALARY NUMERIC, BONUS NUMERIC, PRIMARY KEY (EMPLOYEE_ID, BONUS_YEAR));
```

```
MERGE INTO EMP_BONUS E1
USING (SELECT EMPLOYEE_ID, FIRST_NAME, SALARY, DEPARTMENT_ID
FROM EMPLOYEES) E2 ON (E1.EMPLOYEE_ID = E2.EMPLOYEE_ID) WHEN MATCHED THEN
```

```

UPDATE SET E1.BONUS = E2.SALARY * 0.5
DELETE WHERE (E1.SALARY >= 10000)
WHEN NOT MATCHED THEN
INSERT (E1.EMPLOYEE_ID, E1.BONUS_YEAR, E1.SALARY , E1.BONUS)
VALUES (E2.EMPLOYEE_ID, EXTRACT(YEAR FROM SYSDATE), E2.SALARY,
E2.SALARY * 0.5)
WHERE (E2.SALARY < 10000);

SELECT * FROM EMP_BONUS;

```

EMPLOYEE_ID	BONUS_YEAR	SALARY	BONUS
103	2017	9000	4500
104	2017	6000	3000
105	2017	4800	2400
106	2017	4800	2400
107	2017	4200	2100
111	2017	7700	3850
112	2017	7800	3900
113	2017	6900	3450
115	2017	3100	1550

For more information, see [MERGE](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL doesn't support the MERGE statement. However, it provides two other statements for merging data: REPLACE, and INSERT... ON DUPLICATE KEY UPDATE .

REPLACE deletes a row and inserts a new row if a duplicate key conflict occurs. INSERT... ON DUPLICATE KEY UPDATE performs an in-place update. Both REPLACE and ON DUPLICATE KEY UPDATE rely on an existing primary key and unique constraints. It is not possible to define custom MATCH conditions as with the MERGE statement in Oracle.

REPLACE provides a function similar to INSERT. The difference is that REPLACE first deletes an existing row if a duplicate key violation for a PRIMARY KEY or UNIQUE constraint occurs.

REPLACE is a MySQL extension that is not ANSI compliant. It either performs only an INSERT when no duplicate key violations occur, or it performs a DELETE and then an INSERT if violations occur.

Syntax

```
REPLACE [INTO] <Table Name> (<Column List>) VALUES v(<Values List>)
```

```
REPLACE [INTO] <Table Name> SET <Assignment List: ColumnName = VALUE...>
```

```
REPLACE [INTO] <Table Name> (<Column List>) SELECT ...
```

INSERT ... ON DUPLICATE KEY UPDATE

The `ON DUPLICATE KEY UPDATE` clause of the `INSERT` statement acts as a dual DML hybrid. Similar to `REPLACE`, it executes the assignments in the `SET` clause instead of raising a duplicate key error. `ON DUPLICATE KEY UPDATE` is a MySQL extension that is not ANSI compliant.

```
INSERT [INTO] <Table Name> [<Column List>] VALUES (<Value List>
ON DUPLICATE KEY <Assignment List: ColumnName = Value...>
```

```
INSERT [INTO] <Table Name> SET <Assignment List: ColumnName = Value...>
ON DUPLICATE KEY UPDATE <Assignment List: ColumnName = Value...>
```

```
INSERT [INTO] <Table Name> [<Column List>] SELECT ... ON DUPLICATE KEY
UPDATE <Assignment List: ColumnName = Value...>
```

Migration considerations

Neither `REPLACE` nor `INSERT ... ON DUPLICATE KEY UPDATE` provide a full functional replacement for the `MERGE` statement in Oracle. The key differences are:

- Key violation conditions are mandated by the primary key or unique constraints that exist on the target table. They can't be defined using an explicit predicate.
- There is no alternative for the `WHEN NOT MATCHED BY SOURCE` clause.
- There is no alternative for the `OUTPUT` clause.

The key difference between `REPLACE` and `INSERT ON DUPLICATE KEY UPDATE` is that with `REPLACE`, the violating row is deleted or attempted to be deleted. If foreign keys are in place, the `DELETE` operation may fail, which may fail the entire transaction.

For `INSERT ... ON DUPLICATE KEY UPDATE`, the update is performed on the existing row in place without attempting to delete it.

It should be straightforward to replace most MERGE statements with either REPLACE or INSERT... ON DUPLICATE KEY UPDATE . Alternatively, break down the operations into their constituent INSERT, UPDATE, and DELETE statements.

Examples

Use REPLACE to create a simple one-way, two-table sync.

```
CREATE TABLE SourceTable (Col1 INT NOT NULL PRIMARY KEY,  
    Col2 VARCHAR(20) NOT NULL);  
CREATE TABLE TargetTable (Col1 INT NOT NULL PRIMARY KEY,  
    Col2 VARCHAR(20) NOT NULL);
```

```
INSERT INTO SourceTable (Col1, Col2)  
    VALUES (2, 'Source2'), (3, 'Source3'), (4, 'Source4');  
INSERT INTO TargetTable (Col1, Col2)  
    VALUES (1, 'Target1'), (2, 'Target2'), (3, 'Target3');
```

```
REPLACE INTO TargetTable(Col1, Col2)  
    SELECT Col1, Col2 FROM SourceTable;
```

```
SELECT * FROM TargetTable;
```

For the preceding example, the result looks as shown following.

Col1	Col2
1	Target1
2	Source2
3	Source3
4	Source4

The following example creates a conditional two-way sync using NULL for no change and DELETE from target when not found in source.

```
TRUNCATE TABLE SourceTable;
```

```
INSERT INTO SourceTable(Col1, Col2)  
    VALUES (3, NULL), (4, 'NewSource4'), (5, 'Source5');
```

```
DELETE FROM TargetTable
WHERE Col1 NOT IN (SELECT Col1 FROM SourceTable);
```

```
INSERT INTO TargetTable (Col1, Col2)
SELECT Col1, Col2
FROM SourceTable AS SRC
WHERE SRC.Col1 NOT IN (SELECT Col1 FROM TargetTable);
```

```
UPDATE TargetTable AS TGT
SET Col2 = (SELECT COALESCE(SRC.Col2, TGT.Col2)
FROM SourceTable AS SRC WHERE SRC.Col1 = TGT.Col1)
WHERE TGT.Col1 IN (SELECT Col1 FROM SourceTable);
```

```
SELECT * FROM TargetTable;
```

For the preceding example, the result looks as shown following.

```
Col1 Col2
3     Source3
4     NewSource4
5     Source5
```

Summary

The following table describes similarities, differences, and key migration considerations.

Oracle MERGE feature	Migrate to Aurora MySQL	Comments
Define source set in USING clause.	Define source set in a SELECT query or in a table.	
Define logical duplicate key condition with an ON predicate.	Duplicate key condition mandated by primary key and unique constraints on target table.	
WHEN MATCHED THEN UPDATE	REPLACE or INSERT... ON DUPLICATE KEY UPDATE	When using REPLACE, the violating row is deleted, or

Oracle MERGE feature	Migrate to Aurora MySQL	Comments
		<p>attempted to be deleted.</p> <p>If there are foreign keys in place, the DELETE operation may fail, which may fail the entire transaction. With <code>INSERT ... ON DUPLICATE KEY UPDATE</code>, the update is performed on the existing row in place, without attempting to delete it.</p>
<p>WHEN MATCHED THEN DELETE</p>	<p>DELETE FROM Target WHERE Key IN (SELECT Key FROM Source)</p>	
<p>WHEN NOT MATCHED THEN INSERT</p>	<p>REPLACE or INSERT... ON DUPLICATE KEY UPDATE</p>	<p>When using REPLACE, the violating row is deleted, or attempted to be deleted.</p> <p>If there are foreign keys in place, the DELETE operation may fail, which may fail the entire transaction. With <code>INSERT ... ON DUPLICATE KEY UPDATE</code>, the update is performed on the existing row in place, without attempting to delete it.</p>

For more information, see [REPLACE Statement](#) and [INSERT ... ON DUPLICATE KEY UPDATE Statement](#) in the *MySQL documentation*.

Oracle OLAP functions and MySQL Window functions

The following sections detail the steps for configuring Amazon DMS to handle OLAP functions and Window functions during replication.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	GREATEST and LEAST functions might get different results in MySQL. CONNECT BY isn't supported by MySQL, a workaround is available.

Oracle usage

Oracle OLAP functions extend the functionality of standard SQL analytic functions by providing capabilities to compute aggregate values based on a group of rows. You can apply the OLAP functions to logically partitioned sets of results within the scope of a single query expression. OLAP functions are usually used in combination with Business Intelligence reports and analytics. They can help boost query performance as an alternative to achieving the same result using more complex non-OLAP SQL code.

Common Oracle OLAP functions

Function type	Related functions
Aggregate	average_rank , avg, count, dense_rank , max, min, rank, sum
Analytic	average_rank , avg, count, dense_rank , lag, lag_variance , lead_variance_percent , max, min, rank, row_number , sum, percent_rank , cume_dist , ntile, first_value , last_value

Function type	Related functions
Hierarchical	hier_ancestor , hier_child_count , hier_depth , hier_level , hier_order , hier_parent , hier_top
Lag	lag, lag_variance , lag_variance_percent , lead, lead_variance , lead_variance_percent
OLAP DML	olap_dml_expression
Rank	average_rank , dense_rank , rank, row_number

For more information, see [OLAP Functions](#) and [Functions](#) in the *Oracle documentation*.

MySQL usage

Some Oracle OLAP functions are aggregative functions in Aurora MySQL. For more information, see [Single-Row and Aggregate Functions](#).

You can replace other OLAP functions with window functions, which are currently not available in Aurora MySQL.

Aurora MySQL version 5.7 does not support window functions.

Note

Amazon Relational Database Service (Amazon RDS) for MySQL version 8 supports window functions that for each row from a query perform a calculation using rows related to that row. These include functions such as `RANK()`, `LAG()`, and `NTILE()`. In addition several existing aggregate functions now can be used as window functions such as `SUM()` and `AVG()`. For more information, see [Window Functions](#) in the *MySQL documentation*.

Migration considerations

As a temporary workaround, rewrite the code to remove the use of Window Functions, and revert to using more traditional SQL code solutions.

In most cases, you can find an equivalent SQL query, although it may be less optimal in terms of performance, simplicity, and readability. See the examples below for migrating window functions to code that uses correlated subqueries.

Note

You may want to archive the original code and then reuse it in the future when Aurora MySQL is upgraded to version 8. The documentation for version 8 indicates the window function syntax is ANSI compliant and will be compatible with Oracle PL\SQL syntax.

For more information, see [Window Functions](#) in the *MySQL documentation*.

Examples

The following examples demonstrate ANSI SQL compliant subquery solutions:

Create and populate an OrderItems table.

```
CREATE TABLE OrderItems(  
  OrderID INT NOT NULL, Item VARCHAR(20) NOT NULL,  
  Quantity SMALLINT NOT NULL,  
  PRIMARY KEY(OrderID, Item));  
  
INSERT INTO OrderItems (OrderID, Item, Quantity)  
  VALUES (1, 'M8 Bolt', 100), (2, 'M8 Nut', 100),  
  (3, 'M8 Washer', 200);
```

Rank items based on ordered quantity. This is a workaround for the window ranking function.

```
SELECT Item, Quantity,  
  (SELECT COUNT(*) FROM OrderItems  
   AS OI2 WHERE OI.Quantity > OI2.Quantity) + 1  
  AS QtyRank  
FROM OrderItems AS OI;
```

Calculate the grand total. This is a workaround for a partitioned window aggregate function.

```
SELECT Item, Quantity,
(SELECT SUM(Quantity) FROM OrderItems
 AS OI2 WHERE OI2.OrderID = OI.OrderID)
 AS TotalOrderQty
FROM OrderItems AS OI;
```

For more information, see [Window Function Descriptions](#) in the *MySQL documentation*.

Oracle Transaction Model and MySQL Transactions

Oracle Transaction Model and MySQL Transactions provide mechanisms for grouping SQL statements into logical units of work, ensuring atomicity, consistency, isolation, and durability (ACID properties) during database operations.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Transaction Isolation	In MySQL, the default isolation level is REPEATABLE READ . MySQL doesn't support nested transactions.

Oracle usage

Database transactions are a logical, atomic units of processing containing one or more SQL statements that may run concurrently alongside other transactions. The primary purpose of a transaction is to ensure the ACID model is enforced.

- **Atomicity** — All statements in a transaction are processed as one logical unit, or none are processed. If a single part of a transaction fails, the entire transaction is aborted and no changes are persisted (all or nothing).

- **Consistency** — All data integrity constraints are checked and all triggers are processed before a transaction is processed. If any of the constraints are violated, the entire transaction fails.
- **Isolation** — One transaction isn't affected by the behavior of other concurrent transactions. The effect of a transaction isn't visible to other transactions until the transaction is committed.
- **Durability** — Once a transaction commits, its results will not be lost regardless of subsequent failures. After a transaction completes, changes made by committed transactions are permanent. The database ensures that committed transactions can't be lost.

Database transaction isolation levels

The ANSI/ISO SQL standard (SQL92) defines four levels of isolation. Each level provides a different approach for handling concurrent run of database transactions. Transaction isolation levels manage the visibility of changed data as seen by other running transactions. In addition, when accessing the same data with several concurrent transactions, the selected level of transaction isolation affects the way different transactions interact. For example, if a bank account is shared by two individuals, what will happen if both parties attempt to perform a transaction on the shared account at the same time? One checks the account balance while the other withdraws money. Oracle supports the following isolation levels:

- **Read-uncommitted** — A currently processed transaction can see uncommitted data made by the other transaction. If a rollback is performed, all data is restored to its previous state.
- **Read-committed** — A transaction only sees data changes that were committed. Uncommitted changes("dirty reads") aren't possible.
- **Repeatable read** — A transaction can view changes made by the other transaction only after both transactions issue a COMMIT or both are rolled-back.
- **Serializable** — Any concurrent run of a set of serializable transactions is guaranteed to produce the same effect as running them sequentially in the same order.

Isolation levels affect the following database behavior.

- **Dirty reads** — A transaction can read data that was written by another transaction, but isn't yet committed.
- **Non-repeatable or fuzzy reads** — When reading the same data several times, a transaction can find that the data has been modified by another transaction that has just committed. The same query executed twice can return different values for the same rows.

- **Phantom reads** — Similar to a non-repeatable read, but it is related to new data created by another transaction. The same query run twice can return a different numbers of records.

Isolation level	Dirty reads	Non-repeatable reads	Phantom reads
Read-uncommitted	Permitted	Permitted	Permitted
Read-committed	Not permitted	Permitted	Permitted
Repeatable read	Not permitted	Not permitted	Permitted
Serializable	Not permitted	Not permitted	Not permitted

Oracle isolation levels

Oracle supports the read-committed and serializable isolation levels. It also provides a Read-Only isolation level which isn't a part of the ANSI/ISO SQL standard (SQL92). Read-committed is the default.

- **Read-committed** — Each query that you run within a transaction only sees data that was committed before the query itself. The Oracle database never allows reading dirty pages and uncommitted data. This is the default option.
- **Serializable** — Serializable transactions don't experience non-repeatable reads or phantom reads because they are only able to see changes that were committed at the time the transaction began (in addition to the changes made by the transaction itself performing DML operations).
- **Read-only** — The read-only isolation level doesn't allow any DML operations during the transaction and only sees data committed at the time the transaction began.

Oracle and MySQL Multi-Version Concurrency Control

Oracle uses the Oracle Multiversion Concurrency Controls (MVCC) mechanism to provide automatic read consistency across the entire database and all sessions. Using MVCC, database sessions see data based on a single point in time ensuring only committed changes are viewable. Oracle relies on the System Change Number (SCN) of the current transaction to obtain a consistent view of the

database. Therefore, all database queries only return data committed with respect to the SCN at the time of query run.

Setting isolation levels

Isolation levels can be changed at the transaction and session levels.

Examples

Change the isolation level at the transaction-level.

```
SET TRANSACTION ISOLATION LEVEL READ COMMITTED;  
SET TRANSACTION ISOLATION LEVEL SERIALIZABLE;  
SET TRANSACTION READ ONLY;
```

Change the isolation-level at a session-level.

```
ALTER SESSION SET ISOLATION_LEVEL = SERIALIZABLE;  
ALTER SESSION SET ISOLATION_LEVEL = READ COMMITTED;
```

For more information, see [Transactions](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL supports all four transaction isolation levels described by the SQL:1992 standard: READ UNCOMMITTED, READ COMMITTED, REPEATABLE READ, and SERIALIZABLE.

The default isolation level for Aurora MySQL is REPEATABLE READ. The simplified syntax for setting transaction boundaries in Aurora MySQL is:

```
SET [SESSION] TRANSACTION ISOLATION LEVEL  
[READ WRITE | READ ONLY] |  
REPEATABLE READ | READ COMMITTED |  
READ UNCOMMITTED | SERIALIZABLE]
```

Note

Setting the GLOBAL isolation level is not supported in Aurora MySQL; only session scope can be changed. This behavior is similar to Oracle. Also, the default behavior of

transactions is to use `REPEATABLE READ` and consistent reads. Applications designed to run with `READ COMMITTED` may need to be modified. Alternatively, explicitly change the default to `READ COMMITTED`.

To set the transaction isolation level, configure the `tx_isolation` parameter when using Aurora for MySQL. For more information, see [Oracle Instance Parameters and Aurora MySQL Parameter Groups](#).

In Aurora MySQL, a transaction intent can be optionally specified. Setting a transaction to `READ ONLY` disables the transaction's ability to modify or lock both transactional and non-transactional tables visible to other transactions.

The transaction can still modify or lock temporary tables. This enables internal optimization to improve performance and concurrency. The default is `READ WRITE`.

Note

Amazon Relational Database Service (Amazon RDS) for MySQL version 8, you can use a new `innodb_deadlock_detect` dynamic variable to disable deadlock detection. On high concurrency systems, deadlock detection can cause a slowdown when numerous threads wait for the same lock. At times it may be more efficient to disable deadlock detection and rely on the `innodb_lock_wait_timeout` setting for transaction rollback when a deadlock occurs.

Note

Starting from Amazon RDS for MySQL version 8, InnoDB supports `NOWAIT` and `SKIP LOCKED` options with `SELECT ... FOR SHARE` and `SELECT ... FOR UPDATE` locking read statements. `NOWAIT` causes the statement to return immediately if a requested row is locked by another transaction. `SKIP LOCKED` removes locked rows from the result set. `SELECT ... FOR SHARE` replaces `SELECT ... LOCK IN SHARE MODE` but `LOCK IN SHARE MODE` remains available for backward compatibility. The statements are equivalent. However, `FOR UPDATE` and `FOR SHARE` support `NOWAIT SKIP LOCKED` and `OF tbl_name` options. For more information, see [SELECT Statement](#) in the *MySQL documentation*.

Defining the Beginning of a Transaction

```
START TRANSACTION WITH CONSISTENT SNAPSHOT | READ WRITE | READ ONLY
```

OR

```
BEGIN [WORK]
```

The `WITH CONSISTENT SNAPSHOT` option starts a consistent read Transaction. The effect is the same as issuing a `START TRANSACTION` followed by a `SELECT` from any table. `WITH CONSISTENT SNAPSHOT` doesn't change the transaction isolation level.

A consistent read uses snapshot information to make query results available based on a point-in-time regardless of modifications performed by concurrent transactions. If queried data has been changed by another transaction, the original data is reconstructed using the undo log. This avoids locking issues that may reduce concurrency. With the `REPEATABLE READ` isolation level, the snapshot is based on the time the first read operation is performed. With the `READ COMMITTED` isolation level, the snapshot is reset to the time of each consistent read operation.

Commit work at the end of a transaction:

```
COMMIT [WORK] [AND [NO] CHAIN] [[NO] RELEASE]
```

Roll back work at the end of a transaction:

```
ROLLBACK [WORK] [AND [NO] CHAIN] [[NO] RELEASE]
```

One of the `ROLLBACK` options is `ROLLBACK TO SAVEPOINT <logical_name>`. This command will rollback all changes in current transaction up to the save point mentioned.

Create transaction save point during the transaction

```
SAVEPOINT <logical_name>
```

Note

If the current transaction has a save point with the same name, the old save point is deleted and a new one is set.

The **AND CHAIN** clause causes a new transaction to begin as soon as the current one ends using the same isolation level and access mode as the just-terminated transaction.

The **RELEASE** clause causes the server to disconnect the current session after terminating the current transaction. The **NO** keyword suppresses both **CHAIN** and **RELEASE** completion. This can be useful if the `completion_type` system variable is set to cause chaining or release completion.

Always run with the `autocommit` mode turned on. Set the `autocommit` parameter to 1 on the database side. This is the default value. Also, make sure that the `autocommit` parameter is set to 1 on the application side. This might not be the default value.

Always double-check the `autocommit` settings on the application side. For example, Python drivers such as `MySQLdb` and `PyMySQL` turn off `autocommit` by default.

Aurora MySQL supports auto commit and explicit commit modes. You can change the mode using the system variable `autocommit`, 1 is the default:

```
SET autocommit = {0 | 1}
```

Examples

Run two DML statements within a serializable transaction.

```
SET TRANSACTION ISOLATION LEVEL SERIALIZABLE;
START TRANSACTION;
INSERT INTO Table1 VALUES (1, 'A');

UPDATE Table2 SET Column1 = 'Done' WHERE KeyColumn = 1;
COMMIT;
```

Summary

The following table summarizes the key differences in transaction support and syntax when migrating from Oracle to Aurora MySQL.

Transaction property	Oracle	Aurora MySQL	Comments
Default isolation level	READ COMMITTED	REPEATABLE READ	The Aurora MySQL default isolation level is stricter than

Transaction property	Oracle	Aurora MySQL	Comments
			the Oracle. Evaluate application needs and set appropriately.
Initialize transaction syntax	START TRANSACTION	START TRANSACTION	
Commit transaction	COMMIT [WORK FORCE]	COMMIT [WORK]	If you use only COMMIT or COMMIT WORK, no changes are needed. Otherwise, rewrite FORCE to WORK.
Rollback transaction	ROLLBACK [WORK [TO FORCE]	ROLLBACK [WORK]	If you use only ROLLBACK or ROLLBACK WORK, no changes are needed. Otherwise, rewrite TO and FORCE to WORK.
Set autocommit off or on	SET AUTOCOMMIT ON OFF (SQL*Plus)	SET autocommit = 0 1	
ANSI isolations	REPEATABLE READ READ COMMITTED READ UNCOMMITTED SERIALIZABLE	REPEATABLE READ READ COMMITTED READ UNCOMMITTED SERIALIZABLE	Compatible syntax

Transaction property	Oracle	Aurora MySQL	Comments
MVCC	START TRANSACTION READ COMMITTED	WITH CONSISTENT SNAPSHOT	Aurora MySQL consistent read in READ COMMITTED isolation, is similar to READ COMMITTED in Oracle.
Nested transactions	Supported by starting new transaction or call a procedure or function after transaction start.	Not Supported	Starting a new transaction in Aurora MySQL while another transaction is active causes a COMMIT of the previous transaction.
Transaction chaining	Not Supported	Causes a new transaction to open immediately upon transaction completion.	
Transaction release	Not supported	Causes the client session to disconnect upon transaction completion	

For more information, see [Transaction Isolation Levels](#) in the *MySQL documentation*.

Oracle anonymous block and MySQL transactions or procedures

With Amazon DMS, you can run Oracle anonymous blocks and MySQL transactions or procedures to modify data or perform complex operations during a database migration. An Oracle anonymous block is a set of procedural statements that perform transaction control, data manipulation, or

control flow operations. A MySQL transaction groups multiple SQL statements into a single logical unit of work, while procedures are reusable code objects containing SQL statements.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Different syntax may require code rewrite.

Oracle usage

Oracle PL/SQL is a procedural extension of SQL. The PL/SQL program structure divides the code into blocks distinguished by the following keywords: DECLARE, BEGIN, EXCEPTION, and END.

An unnamed PL/SQL code block (code not stored in the database as a procedure, function, or package) is known as an anonymous block. An anonymous block serves as the basic unit of Oracle PL/SQL and contains the following code sections:

- **The declarative section** (optional) — Contains variables (names, data types, and initial values).
- **The executable section** (mandatory) — Contains executable statements (each block structure must contain at least one executable PL/SQL statement).
- **The exception-handling section** (optional) — Contains elements for handling exceptions or errors in the code.

Examples

Simple structure of an Oracle anonymous block.

```
SET SERVEROUTPUT ON;
BEGIN
DBMS_OUTPUT.PUT_LINE('hello world');
END;
/

hello world
PL/SQL procedure successfully completed.
```

Oracle PL/SQL Anonymous blocks can contain advanced code elements such as functions, cursors, dynamic SQL, and conditional logic. The following anonymous block uses a cursor, conditional logic, and exception-handling.

```
SET SERVEROUTPUT ON;
DECLARE
v_sal_chk          NUMBER;
v_emp_work_years  NUMBER;
v_sql_cmd          VARCHAR2(2000);
BEGIN
FOR v IN (SELECT EMPLOYEE_ID, FIRST_NAME||' '||LAST_NAME AS
EMP_NAME, HIRE_DATE, SALARY FROM EMPLOYEES)
LOOP
v_emp_work_years:=EXTRACT(YEAR FROM SYSDATE) - EXTRACT (YEAR FROM v.hire_date);
IF v_emp_work_years>=10 and v.salary <= 6000 then
DBMS_OUTPUT.PUT_LINE('Consider a Bonus for: '||v.emp_name);
END IF;
END LOOP;
EXCEPTION WHEN OTHERS THEN
DBMS_OUTPUT.PUT_LINE('CODE ERR: '||sqlerrm);
END;
/
```

The preceding example calculates the number of years each employee has worked based on the HIRE_DATE column of the EMPLOYEES table. If the employee has worked for ten or more years and has a salary of \$6000 or less, the system prints the message “Consider a Bonus for: <employee name>”.

For more information, see [Overview of PL/SQL](#) in the *Oracle documentation*.

MySQL usage

You can achieve the similar functionality to Oracle Anonymous Blocks by using the Aurora MySQL START TRANSACTION command or a stored procedure.

For more information, see [Stored Procedures](#) and [Oracle Transaction Model and MySQL Transactions](#).

Conversion functions

With Amazon DMS, you can transform data types between different database platforms during the migration process using conversion functions. Conversion functions define the mapping between data types in the source and target databases, allowing you to handle incompatible data types seamlessly.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	MySQL doesn't support all functions. These unsupported functions require manual creation.

Oracle usage

All databases have their own conversion methods for transforming data between types and performing data manipulation. This section addresses the conversion functions `TO_CHAR` and `TO_NUMBER`.

TO_CHAR

`TO_CHAR` can convert many types of data (mostly number, date, and string) to string. There are many format combinations. Some examples include:

TO_CHAR calls with strings	Results
<code>to_char('0972')</code>	0972
<code>to_char('0972', '9999')</code>	972
<code>to_char('0972', '\$9999.99')</code>	\$972.00
<code>to_char('0972', '\$0009999.99')</code>	\$972.00

TO_CHAR calls with strings	Results
<code>to_char('0972.48', '\$9999.999')</code>	\$972.480

TO_CHAR calls with numbers	Results
<code>to_char(0972)</code>	972
<code>to_char(0972, '9999')</code>	972
<code>to_char(0972, '\$9999.99')</code>	\$972.00
<code>to_char(0972, '\$0009999.99')</code>	\$0000972.00
<code>to_char(0972.48, '\$9999.999')</code>	\$972.480

TO_CHAR calls with date	Results	Description
<code>to_char(sysdate, 'Y YYY')</code>	2013	Year
<code>to_char(sysdate, 'Y Y')</code>	13	Last two digits of the year
<code>to_char(sysdate, 'Y EAR')</code>	TWENTY THIRTEEN	Year in words
<code>to_char(sysdate, 'S YYYY')</code>	2017	S prefixed (-) sign for BC
<code>to_char(sysdate, 'Y ,YYY')</code>	2017	Year with comma
<code>to_char(sysdate, 'M ONTH')</code>	SEPTEMBER	Complete month

TO_CHAR calls with date	Results	Description
<code>to_char(sysdate, 'MON')</code>	SEP	Three-letter month format
<code>to_char(sysdate, 'MM')</code>	9	Month of the year
<code>to_char(sysdate, 'WW')</code>	4	Week of the current month
<code>to_char(sysdate, 'WW')</code>	36	Week of the year (1 - 53)
<code>to_char(sysdate, 'DAY')</code>	SATURDAY	Name of the day
<code>to_char(sysdate, 'DD')</code>	30	Day in number format
<code>to_char(sysdate, 'D')</code>	7	Day of the week (1 - 7)
<code>to_char(sysdate, 'DDD')</code>	273	Day of the year (1 - 366)
<code>to_char(sysdate, 'DY')</code>	SAT	Short form of the day
<code>to_char(sysdate, 'HH')</code>	9	Hour (1 - 12)
<code>to_char(sysdate, 'HH12')</code>	9	Hour in 12 hours format
<code>to_char(sysdate, 'HH24')</code>	21	Hour in 24 hours format
<code>to_char(sysdate, 'MI')</code>	15	Minute (0 - 59)

TO_CHAR calls with date	Results	Description
<code>to_char(sysdate, 'S')</code>	24	Second (0 - 59)
<code>to_char(sysdate, 'SSSS')</code>	79100	Seconds after midnight (0 - 86399)
<code>to_char(sysdate, 'PM')</code>	PM	AM or PM
<code>to_char(sysdate, 'AM')</code>	PM	AM or PM
<code>to_char(sysdate, 'DL')</code>	Saturday, February 23, 2017	Date in long format
<code>to_char(sysdate, 'Q')</code>	3	Quarter of the Year (1 - 4)

TO_NUMBER

TO_NUMBER converts one of the following to number data types: CHAR, VARCHAR2, NCHAR, NVARCHAR2, BINARY_FLOAT, or BINARY_DOUBLE. When converting one of the first four types, you can use the format parameter for the returned number.

The format parameter specifies one of the following options:

Example data to convert	Format parameter	Results
-1234567890	9999999999S	'1234567890-'
0	99.99	' .00'
0.1	99.99	' .10'
-0.2	99.99	' -.20'
0	9999	' 0'

MySQL usage

For more information, see [Single-Row and Aggregate Functions](#).

Oracle and MySQL cursors

With Amazon DMS, you can efficiently migrate data from Oracle and MySQL databases to other database services or engines, including handling complex data types, such as cursors. A cursor is a database object that allows traversal over rows from a query, facilitating operations like looping through result sets.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Cursors	Minor differences in syntax may require some code rewrite. MySQL doesn't support %ISOPEN, %ROWTYPE, and %BULK_ROWCOUNT .

Oracle usage

PL/SQL cursors are pointers to data sets on which application logic can iterate. The data sets hold rows returned by SQL statements. You can refer to the active data set in named cursors from within a program.

There are two types of PL/SQL cursors:

- **Implicit cursors** are session cursors constructed and managed by PL/SQL automatically without being created or defined by a user. PL/SQL opens an implicit cursor each time you run a SELECT or DML statement. Implicit cursors are also called SQL cursors.
- **Explicit cursors** are session cursors created, constructed, and managed by a user. Cursors are declared and defined by naming it and associating it with a query. Unlike an implicit cursor, you can reference an explicit cursor using its name. An explicit cursor is called a named cursor.

Examples

The following examples demonstrate cursor usage:

1. Define an explicit PL/SQL cursor named c1.
2. The cursor runs an SQL statement to return rows from the database.
3. The PL/SQL loop reads data from the cursor, row by row, and stores the values into two variables: v_lastname and v_jobid.
4. The loop uses the %NOTFOUND attribute to terminate when the last row is read from the database.

```
DECLARE
  CURSOR c1 IS
    SELECT last_name, job_id FROM employees
    WHERE REGEXP_LIKE (job_id, 'S[HT]_CLERK')
    ORDER BY last_name;
  v_lastname employees.last_name%TYPE; -- variable to store last_name
  v_jobid employees.job_id%TYPE; -- variable to store job_id
BEGIN
  OPEN c1;
  LOOP -- Fetches 2 columns into variables
    FETCH c1 INTO v_lastname, v_jobid;
    EXIT WHEN c1%NOTFOUND;
  END LOOP;
  CLOSE c1;
END;
```

1. Define an implicit PL/SQL cursor using a FOR Loop.
2. The cursor runs a query and stores values returned into a record.
3. A loop iterates over the cursor data set and prints the result.

```
BEGIN
  FOR item IN
    (SELECT last_name, job_id FROM employees WHERE job_id LIKE '%MANAGER%'
     AND manager_id > 400 ORDER BY last_name) LOOP
    DBMS_OUTPUT.PUT_LINE('Name = ' || item.last_name || ', Job = ' || item.job_id);
  END LOOP;
END;
```

/

For more information, see [Explicit Cursor Declaration and Definition](#) and [Implicit Cursor Attribute](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL supports cursors only within stored routines, functions and stored procedures.

Unlike Oracle, which offers an array of cursor types, Aurora MySQL Cursors have the following characteristics:

- **Not sensitive** — The server can choose to either make a copy of its result table or to access the source data as the cursor progresses.
- **Read-only** — Cursors can't be updated.
- **Not scrollable** — Cursors can only be traversed in one direction and cannot skip rows. The only supported cursor advance operation is `FETCH NEXT`.

Cursor declarations must appear before handler declarations and after variable and condition declarations.

Similar to Oracle, cursors are declared with the `DECLARE CURSOR`, opened with `OPEN`, fetched with `FETCH`, and closed with `CLOSE`.

Declare Cursor

```
DECLARE <Cursor Name> CURSOR FOR <Cursor SELECT Statement>
```

The `DECLARE CURSOR` statement instantiates a cursor object and associates it with a `SELECT` statement. This `SELECT` is then used to retrieve the cursor rows.

To fetch the rows, use the `FETCH` statement. As mentioned before, only `FETCH NEXT` is supported. The number of output variables specified in the `FETCH` statement must match the number of columns retrieved by the cursor.

Aurora MySQL cursors have additional characteristics:

- `SELECT INTO` is not allowed in a cursor.

- Stored routines can have multiple cursor declarations, but all cursors declared in a given code block must have a unique name.
- Cursors can be nested.

Open cursor

```
OPEN <Cursor Name>;
```

The OPEN command populates the cursor with the data, either dynamically or in a temporary table, and reads the first row for consumption by the FETCH statement.

Fetch cursor

```
FETCH [[NEXT] FROM] <Cursor Name> INTO <Variable 1> [, <Variable n>]
```

The FETCH statement retrieves the current pointer row, assigns the column values to the variables listed in the FETCH statement, and advances the cursor pointer by one row. If the row is not available, meaning the cursor has been exhausted, a No Data condition is raised with an SQLSTATE value of '020000'. To catch this condition, or the alternative NOT FOUND condition, you must create a condition handler.

Note

Carefully plan your error handling flow. The same condition might be raised by SELECT statements or cursors other than the one you intended. Place operations within BEGIN ... END blocks to associate each cursor with its own handler.

Close cursor

```
CLOSE <Cursor Name>;
```

The CLOSE statement closes an open cursor. If the cursor with the specified name does not exist, an error is raised. If a cursor is not explicitly closed, Aurora MySQL closes it automatically at the end of the BEGIN ... END block in which it was declared.

Examples

The following example uses a cursor to iterate over source rows and merges into an OrderItems table.

Create an OrderItems table.

```
CREATE TABLE OrderItems(OrderID INT NOT NULL,  
    Item VARCHAR(20) NOT NULL,  
    Quantity SMALLINT NOT NULL,  
    PRIMARY KEY(OrderID, Item));
```

Create and populate the SourceTable.

```
CREATE TABLE SourceTable (  
    OrderID INT,  
    Item VARCHAR(20),  
    Quantity SMALLINT,  
    PRIMARY KEY (OrderID, Item));  
  
INSERT INTO SourceTable (  
    OrderID, Item, Quantity)  
VALUES (1, 'M8 Bolt', 100),  
    (2, 'M8 Nut', 100),  
    (3, 'M8 Washer', 200);
```

Create a procedure to loop through SourceTable and insert rows.

```
CREATE PROCEDURE LoopItems()  
BEGIN  
    DECLARE done INT DEFAULT FALSE;  
    DECLARE var_OrderID INT;  
    DECLARE var_Item VARCHAR(20);  
    DECLARE var_Quantity SMALLINT;  
    DECLARE ItemCursor CURSOR FOR SELECT OrderID, Item, Quantity FROM SourceTable;  
    DECLARE CONTINUE HANDLER FOR NOT FOUND SET done = TRUE;  
    OPEN ItemCursor;  
    CursorStart: LOOP  
    FETCH NEXT FROM ItemCursor INTO var_OrderID, var_Item, var_Quantity;  
    IF Done THEN LEAVE CursorStart;  
    END IF;  
    INSERT INTO OrderItems (OrderID, Item, Quantity)
```

```
VALUES (var_OrderID, var_Item, var_Quantity);
END LOOP;
CLOSE ItemCursor;
END;
```

Run the stored procedure.

```
CALL LoopItems();
```

Select all rows from the OrderItems table.

```
SELECT * FROM OrderItems;
```

OrderID	Item	Quantity
1	M8 Bolt	100
2	M8 Nut	100
3	M8 Washer	200

Summary

Action	Oracle	Aurora MySQL
Declare a bound explicit cursor	<pre>CURSOR c1 IS SELECT * FROM employees ;</pre>	<pre>DECLARE c1 CURSOR FOR SELECT * FROM employees;</pre>
Open a cursor	<pre>OPEN c1;</pre>	<pre>OPEN c1;</pre>
Move the cursor to the next row and fetch into a record variable (rowvar was declared in the DECLARE section)	<pre>FETCH c1 INTO rowvar;</pre>	<pre>FETCH NEXT FROM c1 INTO rowvar;</pre>
Move the cursor to the next row and fetch into multiple scalar data types (emp_id, emp_name, and salary were	<pre>FETCH c1 INTO emp_id, emp_name, salary;</pre>	<pre>FETCH NEXT FROM c1 INTO emp_id, emp_name, salary;</pre>

Action	Oracle	Aurora MySQL
declared in the DECLARE section)		
Iterate through an implicit cursor using a loop	<pre>FOR item IN (SELECT last_name, job_id FROM employees WHERE job_id LIKE '%CLERK%' AND manager_id > 120 ORDER BY last_name) LOOP << do something >> END LOOP;</pre>	N/A
Declare a cursor with variables	<pre>CURSOR c1 (key NUMBER) IS SELECT * FROM employees WHERE id = key;</pre>	<pre>SET @sqltext1 := CONCAT ('DECLARE c1 CURSOR FOR SELECT * FROM employees WHERE id =',key); PREPARE stmt1 FROM @sqltext1; EXECUTE stmt1;</pre>
Open a cursor with variables	<pre>OPEN c1(2);</pre>	Use regular OPEN after declaring the CURSOR using EXECUTE and PREPARE with variables
Exit a loop after no data found	<pre>EXIT WHEN c1%NOTFOUND;</pre>	<pre>DECLARE CONTINUE HANDLER FOR NOT FOUND SET done = TRUE; And in the fetching loop insert: IF done THEN LEAVE;</pre>

Action	Oracle	Aurora MySQL
Detect if a cursor has rows remaining in its dataset	%FOUND	N/A
Determine how many rows were affected from any DML statement	%BULK_ROWCOUNT	Use counters
Determine which DML run failed with the relevant error code	%BULK_EXCEPTIONS	N/A
Detect if the cursor is open	%ISOPEN	N/A
Detect if the cursor has no rows remaining in its dataset	%NOTFOUND	N/A
Return the number of rows affected by a cursor	%ROWCOUNT	N/A

For more information, see [Cursors](#) in the *MySQL documentation*.

Oracle DBMS_DATAPUMP and MySQL integration with Amazon S3

With Amazon DMS, you can migrate data from Oracle databases to Amazon S3 using Oracle DBMS_DATAPUMP, and load data from Amazon S3 into MySQL-compatible databases. Oracle DBMS_DATAPUMP provides a way to transfer data objects between Oracle databases or export them to an operating system file. MySQL integration with Amazon S3 lets you use an Amazon S3 bucket as a data source or destination for loading and unloading data.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	No equivalent tool

Oracle usage

The DBMS_DATAPUMP package provides Oracle Data Pump functionality that can be run within the database.

The DBMS_DATAPUMP package subprograms are:

- ADD_FILE — Adds a relevant file to the dump file set.
- ATTACH — Connects the DATAPUMP job.
- DATA_FILTER — Filters rows.
- DETACH — Disconnects from a DATAPUMP operation.
- GET_DUMPFILE_INFO — Retrieves information about a specified dump file.
- GET_STATUS — Retrieves status of the running DATAPUMP operation.
- LOG_ENTRY — Writes a message into the log file.
- METADATA_FILTER — Filters the items to be include in the operation.
- METADATA_REMAP — Remaps the object to new names.
- METADATA_TRANSFORM — Specifies transformations to be applied to objects.
- OPEN — Declares a new job.
- SET_PARALLEL — Set the parallelism of the job.
- SET_PARAMETER — Specifies job processing options.
- START_JOB — Runs a job.
- STOP_JOB — Terminates a job.
- WAIT_FOR_JOB — Runs a job until it either completes normally or stops.

Examples

The following example shows how to export the HR schema. It assumes all directories have already been created and the user has all required privileges.

```
DECLARE
loopidx NUMBER;
job_handle NUMBER;
percent_done NUMBER;
job_state VARCHAR2(30);
err ku$_LogEntry;
job_status ku$_JobStatus;
job_desc ku$_JobDesc;
obj_stat ku$_Status;
BEGIN

job_handle := DBMS_DATAPUMP.OPEN('EXPORT', 'SCHEMA', NULL, 'EXP_SAMP', 'LATEST');

DBMS_DATAPUMP.ADD_FILE(job_handle, 'hr.dmp', 'DMPDIR');

DBMS_DATAPUMP.METADATA_FILTER(job_handle, 'SCHEMA_EXPR', 'IN (''HR'')');

DBMS_DATAPUMP.START_JOB(job_handle);

percent_done := 0;
job_state := 'UNDEFINED';
while (job_state != 'COMPLETED') and (job_state != 'STOPPED') loop
dbms_datapump.get_status(job_handle,
dbms_datapump.ku$_status_job_error +
dbms_datapump.ku$_status_job_status +
dbms_datapump.ku$_status_wip, -1, job_state, obj_stat);
job_status := obj_stat.job_status;

/* HERE YOU CAN PRINT THE STATUS */

if job_status.percent_done != percent_done then
percent_done := job_status.percent_done;
end if;

if (bitand(obj_stat.mask, dbms_datapump.ku$_status_wip) != 0) then
err := obj_stat.wip;
else
if (bitand(obj_stat.mask, dbms_datapump.ku$_status_job_error) != 0)
```

```

then
  err := obj_stat.error;
else
  err := null;
end if;
end if;

if err is not null then
  loopidx := err.FIRST;
  while loopidx is not null loop
    loopidx := err.NEXT(loopidx);
  end loop;
end if;
end loop;

dbms_datapump.detach(job_handle);
END;
/

```

For more information, see [Overview of Oracle Data Pump](#) in the *Oracle documentation*.

MySQL usage

There is no feature in MySQL fully equivalent to the Oracle DBMS_DATAPUMP package, but there are tools and features that achieve the same functionality.

To export data from the database to the file system, use the `SELECT INTO OUTFILE S3` command. To import data from the filesystem, use the `LOAD DATA FROM S3` command.

To achieve the most functionality, this feature can be mixed with metadata tables and events to handle the operations.

For more information, see [Oracle External Tables and MySQL Integration with Amazon S3](#).

Summary

Feature	Oracle DBMS_DATAPUMP	Aurora integration with S3
Add a relevant file to the dump file set	ADD_FILE	Use metadata table
Connect the DATAPUMP job	ATTACH	Query session status

Feature	Oracle DBMS_DATAPUMP	Aurora integration with S3
Filter rows to be handled	DATA_FILTER	Use WHERE clause in your SELECT
Disconnect from DATAPUMP operation	DETACH	Not required
Retrieve information about a specified dump file	GET_DUMPFILINFO	Use metadata table
Retrieve the status of the running DATAPUMP operation	GET_STATUS	Query session status
Write a message into the log file	LOG_ENTRY	Write to metadata tables
Filter the items included in the operation	METADATA_FILTER	Export the objects
Remap the object to new names	METADATA_REMAP	LOAD DATA INTO different table name
Specified transformations to be applied to objects	METADATA_TRANSFORM	Not required
Declare a new job	OPEN	Use LOAD DATA or SAVE OUTFILE
Set the parallelism of the job	SET_PARALLEL	Use parallel in your SELECT
Specify job-processing options	SET_PARAMETER	Not required
Run a job	START_JOB	Use LOAD DATA or SAVE OUTFILE
Terminate a job	STOP_JOB	Kill session

Feature	Oracle DBMS_DATAPUMP	Aurora integration with S3
Run a job until it either completes normally or stops	WAIT_FOR_JOB	Use LOAD DATA or SAVE OUTFILE

Oracle DBMS_OUTPUT and MySQL SELECT

Oracle DBMS_OUTPUT is a package that lets you send messages from stored procedures, functions, and anonymous blocks to a message buffer. MySQL SELECT is a statement used to retrieve data from one or more tables in a MySQL database. The following sections will provide details on using DBMS_OUTPUT in Oracle and SELECT statements in MySQL with Amazon DMS.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		DBMS_OUTPUT	Different paradigm and syntax requires application and drivers rewrite.

Oracle usage

The Oracle DBMS_OUTPUT package is typically used for debugging or for displaying output messages from PL/SQL procedures.

Examples

In the following example, DBMS_OUTPUT with PUT_LINE is used with a combination of bind variables to dynamically construct a string and print a notification to the screen from within an Oracle PL/SQL procedure. In order to display notifications on to the screen, you must configure the session with SET SERVEROUTPUT ON.

```
SET SERVEROUTPUT ON
DECLARE
CURSOR c1 IS
```

```
SELECT last_name, job_id FROM employees
WHERE REGEXP_LIKE (job_id, 'S[HT]_CLERK')
ORDER BY last_name;
v_lastname employees.last_name%TYPE; -- variable to store last_name
v_jobid employees.job_id%TYPE; -- variable to store job_id
BEGIN
OPEN c1;
LOOP -- Fetches 2 columns into variables
FETCH c1 INTO v_lastname, v_jobid;
DBMS_OUTPUT.PUT_LINE ('The employee id is:' || v_jobid || ' and his last name is:' ||
v_lastname);
EXIT WHEN c1%NOTFOUND;
END LOOP;
CLOSE c1;
END;
```

In addition to the output of information on the screen, the PUT and PUT_LINE procedures in the DBMS_OUTPUT package enable you to place information in a buffer that can be read later by another PL/SQL procedure or package. You can display the previously buffered information using the GET_LINE and GET_LINES procedures.

For more information, see [DBMS_OUTPUT](#) in the *Oracle documentation*.

MySQL usage

You can use SELECT to display output messages in Aurora MySQL.

Examples

```
delimiter //

CREATE PROCEDURE emp_counter (param1 INTEGER)
BEGIN
SELECT "" 'OUTPUT: Before count';
SELECT COUNT(*) INTO param1 FROM EMPS;
SELECT concat('Employees count: ', param1) as '';
SELECT "" 'OUTPUT: After count';
END//

delimiter ;
call simpleproc1(1);
```

```

OUTPUT: Before count
1 row in set (0.19 sec)

Employees count: 8
1 row in set (0.20 sec)

OUTPUT: After count
1 row in set (0.21 sec)

Query OK, 0 rows affected (0.22 sec)

```

Note

Use double quotation marks with SELECT for cleaner display. Otherwise, messages are displayed twice, both as header and value.

For more information, see [SELECT Statement](#) in the *MySQL documentation*.

Oracle DBMS_RANDOM and MySQL RAND function

With Amazon DMS, you can generate random numbers or values during data migration from Oracle to MySQL or vice versa. Oracle's DBMS_RANDOM package and MySQL's RAND function provide methods for generating random data, which can be useful for tasks like creating test data, simulating real-world scenarios, or introducing randomness into algorithms.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	Different syntax and missing options may require code rewrite.

Oracle usage

Oracle DBMS_RANDOM package provides functionality for generating random numbers or strings as part of an SQL statement or PL/SQL procedure.

The DBMS_RANDOM Package stored procedures include:

- **NORMAL** — Returns random numbers in a standard normal distribution.
- **SEED** — Resets the seed that generates random numbers or strings.
- **STRING** — Returns a random string.
- **VALUE** — Returns a number greater than or equal to 0 and less than 1 with 38 digits to the right of the decimal. Alternatively, you could generate a random number greater than or equal to a low parameter and less than a high parameter.

DBMS_RANDOM.RANDOM produces integers in the range $[-2^{31}, 2^{31}]$.

DBMS_RANDOM.VALUE produces numbers in the range $[0,1]$ with 38 digits of precision.

Examples

Generate a random number.

```
select dbms_random.value() from dual;

DBMS_RANDOM.VALUE()
.859251508

select dbms_random.value() from dual;

DBMS_RANDOM.VALUE()
.364792387
```

Generate a random string. The first character determines the returned string type and the number specifies the length.

```
select dbms_random.string('p',10) from dual;
DBMS_RANDOM.STRING('P',10)

la'?[Q&/2

select dbms_random.string('p',10) from dual;
DBMS_RANDOM.STRING('P',10)

t?!Gf2M60q
```

For more information, see [DBMS_RANDOM](#) in the *Oracle documentation*.

MySQL usage

The MySQL RAND function is not fully equivalent to Oracle DBMS_RANDOM because it does not generate string values. However, there are other functions in that can be used in combination to achieve full functionality.

RAND function returns a random floating-point value v in the range $0 \leq v < 1.0$.

You can use the RAND function with a seed value to reset the seed. If an integer argument N is specified, it is used as the seed value:

- With a constant initializer argument, the seed is initialized once when the statement is prepared and prior to execution.
- With a non-constant initializer argument such as a column name, the seed is initialized with the value for each invocation of RAND().

Examples

Generate a random number.

```
select RAND();

RAND()
0.30244802525494996
```

To obtain a random integer R in the range $i \leq R < j$, use the expression $\text{FLOOR}(i + \text{RAND}() * (j - i))$. For example, to obtain a random integer in the range $7 \leq R < 12$, use:

```
SELECT FLOOR(7 + (RAND() * 5));

FLOOR(7 + (RAND() * 5))
8
```

Generate an eight-character string of digits.

```
SELECT SUBSTRING(MD5(RAND())) FROM 1 FOR 8);
```

Generate an eight-character string containing characters only.

```
SELECT concat(substring('ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxy',
  rand()*52+1, 1),
substring('ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxy', rand()*52+1, 1))
```

For more information, see [RAND\(\)](#) in the *MySQL documentation*.

Oracle DBMS_REDEFINITION and MySQL tables and triggers

The following sections provide detailed guidance on leveraging the Oracle DBMS_REDEFINITION and MySQL tables and triggers features during database migration using Amazon DMS.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	MySQL doesn't support DBMS_REDEFINITION .

Oracle usage

The Oracle DBMS_REDEFINITION package can be used to reorganize tables while they perform DML operations. Use this package to reclaim space due to a high watermark or to change the table's DDL.

Oracle uses materialized views to track changes on the master table and then applies those changes in refresh synchronization.

Examples

Run online redefinition.

- `DBMS_REDEFINITION.CAN_REDEF_TABLE` — Determines if the table can be redefined.

- `DBMS_REDEFINITION.START_REDEF_TABLE` — Start the online redefinition.
- `DBMS_REDEFINITION.SYNC_INTERIM_TABLE` — Synchronize tables with interim data.
- `DBMS_REDEFINITION.FINISH_REDEF_TABLE` — Complete redefinition.

```
EXEC DBMS_REDEFINITION.CAN_REDEF_TABLE('HR', 'EMPLOYEES');
CREATE TABLE employees2 AS SELECT * FROM employees WHERE 1=2;

EXEC DBMS_REDEFINITION.START_REDEF_TABLE
  ('HR', 'EMPLOYEES', 'EMPLOYEES2', '*');

EXEC DBMS_REDEFINITION.SYNC_INTERIM_TABLE
  ('HR', 'EMPLOYEES', 'EMPLOYEES2');

ALTER TABLE employees2 ADD
  (CONSTRAINT emp_pk2 PRIMARY KEY (empno) USING INDEX);

EXEC DBMS_REDEFINITION.FINISH_REDEF_TABLE
  ('HR', 'EMPLOYEES', 'EMPLOYEES2');

DROP TABLE employees2;
```

For more information, see [DBMS_REDEFINITION](#) in the *Oracle documentation*.

MySQL usage

MySQL has no equivalent to Oracle for automatically rebuilding tables or syncing between two tables. However, you can sync data from one table to another using `CREATE TABLE AS SELECT` or `mysqldump`. After the table is copied, the delta rows can be copied using triggers. Once the application is ready to use the new table, it is synced.

If a table has sequence columns, the last value in the sequence is retained when the table is copied.

For more information, see [Trigger Syntax and Examples](#), [CREATE TABLE ... SELECT Statement](#), and [mysqldump — A Database Backup Program](#) in the *MySQL documentation*.

Oracle DBMS_SQL

When working with Oracle databases migrated to Amazon you can use `DBMS_SQL` to maintain application functionality to run dynamic SQL statements. You can also use `DBMS_SQL` for

automating database operations. The following sections cover the details of using DBMS_SQL with code examples.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	Different paradigm and syntax will require application and drivers rewrite.

Oracle usage

The DBMS_SQL package provides an interface to parse and run dynamic SQL statements, DML commands, and DDL commands (usually from within a PL/SQL package, function, or procedure). DBMS_SQL enables very granular control of SQL cursors and can improve cursor performance in certain cases.

Examples

The following examples demonstrates how to manually open, parse, bind, run, and fetch data from a cursor using the DBMS_SQL PL/SQL interface.

1. Use DBMS_SQL.OPEN_CURSOR to open a blank cursor and return the cursor handle.
2. Use DBMS_SQL.PARSE to parse the statement into the referenced cursor.
3. Use DBMS_SQL.BIND_VARIABLES to attach the value for the bind variable with the cursor.
4. Use DBMS_SQL.EXECUTE to run the cursor.
5. Use DBMS_SQL.GET_NEXT_RESULT to iterate over the cursor, fetching the next result.
6. Use DBMS_SQL.CLOSE_CURSOR to close the cursor.

```
DECLARE
c1          INTEGER;
rc1        SYS_REFCURSOR;
n          NUMBER;
```

```
first_name  VARCHAR2(50);
last_name   VARCHAR2(50);
email       VARCHAR2(50);
phone_number VARCHAR2(50);
job_title   VARCHAR2(50);
start_date  DATE;
end_date    DATE;
BEGIN
c1 := DBMS_SQL.OPEN_CURSOR(TRUE);
DBMS_SQL.PARSE
  (c1, 'BEGIN emp_info(:id); END;', DBMS_SQL.NATIVE);
DBMS_SQL.BIND_VARIABLE(c1, ':id', 176);
n := DBMS_SQL.EXECUTE(c1);
-- Get employee info
DBMS_SQL.GET_NEXT_RESULT(c1, rc1);
FETCH rc1 INTO first_name, last_name, email, phone_number;
-- Get employee job history
DBMS_SQL.GET_NEXT_RESULT(c1, rc1);
LOOP
FETCH rc1 INTO job_title, start_date, end_date;
EXIT WHEN rc1%NOTFOUND;
END LOOP;
DBMS_SQL.CLOSE_CURSOR(c1);
END;
/
```

The `DBMS_SQL` package includes three other procedures.

- `RETURN_RESULT` — Gets a result set and returns it to the client. Because the procedure already returns a result set, the invoker doesn't have to know the format of the result or the columns it contains. This option is new in Oracle 12c and is most often used with SQL*Plus.
- `TO_REFCURSOR` — When using `DBMS_SQL.OPEN_CURSOR`, the numeric cursor ID is returned. If you know the structure of the result of the cursor, you can call the `TO_REFCURSOR` procedure, stop working with `DBMS_SQL`, and move to regular commands such as `FETCH`, `WHEN CURSOR %not found`, and others. Before using `TO_REFCURSOR`, use the procedures `OPEN_CURSOR`, `PARSE`, and `EXECUTE`.
- `TO_CURSOR_NUMBER` — Gets a cursor opened in native dynamic SQL. After the cursor is open, it can be converted to a number or cursor id and then managed using `DBMS_SQL` procedures.

For more information, see [DBMS_SQL](#) in the *Oracle documentation*.

MySQL usage

There is no equivalent feature for the DBMS_SQL package in MySQL. The only options for Aurora MySQL are:

- Procedures or functions.
- Prepare and run statements.

For more information, see [CREATE PROCEDURE and CREATE FUNCTION Statements](#) in the *MySQL documentation*.

Oracle EXECUTE IMMEDIATE and MySQL EXECUTE and PREPARE statements

With Amazon DMS, you can migrate databases between different database platforms, including Oracle and MySQL, by leveraging features, such as Oracle's EXECUTE IMMEDIATE statement and MySQL's EXECUTE and PREPARE statements.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		EXECUTE IMMEDIATE	Make sure that you use the PREPARE command in MySQL. MySQL doesn't support running SQL with results and bind variables or anonymous blocks using EXECUTE.

Oracle usage

You can use Oracle EXECUTE IMMEDIATE statement to parse and run a dynamic SQL statement or an anonymous PL/SQL block. It also supports bind variables.

Examples

Run a dynamic SQL statement from within a PL/SQL procedure:

1. Create a PL/SQL procedure named `raise_sal`.
2. Define a SQL statement with a dynamic value for the column name included in the WHERE statement.
3. Use the EXECUTE IMMEDIATE command supplying the two bind variables to be used as part of the SELECT statement: `amount` and `col_val`.

```
CREATE OR REPLACE PROCEDURE raise_sal (col_val NUMBER,
emp_col VARCHAR2, amount NUMBER) IS
  col_name VARCHAR2(30);
  sql_stmt VARCHAR2(350);
BEGIN
  -- determine if a valid column name has been given as input
  SELECT COLUMN_NAME INTO col_name FROM USER_TAB_COLS
  WHERE TABLE_NAME = 'EMPLOYEES' AND COLUMN_NAME = emp_col;

  -- define the SQL statement (with bind variables)
  sql_stmt := 'UPDATE employees SET salary = salary + :1 WHERE ' ||
  col_name || ' = :2';

  -- Run the command
  EXECUTE IMMEDIATE sql_stmt USING amount, col_val;
END raise_sal;
/
```

4. Run the DDL operation from within an EXECUTE IMMEDIATE command.

```
EXECUTE IMMEDIATE 'CREATE TABLE link_emp (idemp1 NUMBER, idemp2 NUMBER)';
EXECUTE IMMEDIATE 'ALTER SESSION SET SQL_TRACE TRUE';
```

5. Run an anonymous block with bind variables using EXECUTE IMMEDIATE.

```
EXECUTE IMMEDIATE 'BEGIN raise_sal (:col_val, :col_name, :amount); END;'
  USING 134, 'EMPLOYEE_ID', 10;
```

For more information, see [EXECUTE IMMEDIATE Statement](#) in the *Oracle documentation*.

MySQL usage

The EXECUTE command in MySQL runs commands that were prepared by the PREPARE command. It can also run DDL statements and retrieve data using SQL commands. Similar to Oracle, you can use the MySQL EXECUTE command with bind variables.

The PREPARE command can receive a SELECT, INSERT, UPDATE, DELETE, or VALUES statement and parse it with a user-specified qualifying name so that you can use the EXECUTE command later without the need to re-parse the SQL statement for each run.

- Statement names are not case-sensitive. A Statement name is either a string literal or a user variable containing the text of the SQL statement.
- If a prepared statement with the given name already exists, it is deallocated implicitly before the new statement is prepared.
- The scope of a prepared statement is the session in which it is created.

Examples

Run a SQL SELECT query with the table name as a dynamic variable using bind variables. This query returns the number of employees under a manager with a specific ID.

```
PREPARE stmt1 FROM 'SELECT count(*) FROM employees WHERE ID=?';
SET @man_id = 3;
EXECUTE stmt1 USING @a;

count(*)
2
```

Run a DML command with no variables and then with variables.

```
PREPARE stmt1 FROM 'INSERT INTO numbers (a) VALUES (1)';
EXECUTE stmt1;

PREPARE stmt1 FROM 'INSERT INTO numbers (a) VALUES (?)';
SET @man_id = 3;
EXECUTE stmt1 USING @a;
```

Run a DDL command.

```
PREPARE stmt1 FROM 'CREATE TABLE numbers (num integer)';
EXECUTE stmt1;
```

Summary

Functionality	Oracle EXECUTE IMMEDIATE	MySQL EXECUTE and PREPARE
Run SQL with results and bind variables	<pre>EXECUTE IMMEDIATE 'select salary from employees WHERE ' col_name ' = :1' INTO amount USING col_val;</pre>	N/A
Run DML with variables and bind variables	<pre>EXECUTE IMMEDIATE 'UPDATE employees SET salary = salary + :1 WHERE ' col_name ' = :2' USING amount, col_val;</pre>	<pre>PREPARE stmt1 FROM 'UPDATE employees SET salary = salary + ? WHERE ? = ?' EXECUTE stmt1 USING @amount,@ col,@colval;</pre>
Run DDL	<pre>EXECUTE IMMEDIATE 'CREATE TABLE link_emp (idemp1 NUMBER, idemp2 NUMBER)';</pre>	<pre>PREPARE stmt1 FROM 'CREATE TABLE link_emp (idemp1 INTEGER, idemp2 INTEGER)' EXECUTE stmt1;</pre>
Run an anonymous block	<pre>EXECUTE IMMEDIATE 'BEGIN DBMS_OUTPUT.PUT_LINE ("Anonymous Block"); END;';</pre>	N/A

For more information, see [EXECUTE Statement](#) and [PREPARE Statement](#) in the *MySQL documentation*.

Oracle procedures and functions and MySQL stored procedures

By migrating procedures, functions, and stored procedures, you can preserve existing business logic and functionality in the new database. The following sections provide detailed steps for migrating these database objects using Amazon DMS, ensuring a smooth transition while maintaining data integrity and application compatibility.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Stored Procedures	Syntax and option differences.

Oracle usage

PL/SQL is Oracle built-in database programming language providing several methods to store and run reusable business logic from within the database. Procedures and functions are reusable snippets of code created using the `CREATE PROCEDURE` and the `CREATE FUNCTION` statements.

Stored procedures and stored functions are PL/SQL units of code consisting of SQL and PL/SQL statements that solve specific problems or perform a set of related tasks.

Procedure is used to perform database actions with PL/SQL.

Function is used to perform a calculation and return a result.

Privileges for creating procedures and functions

To create procedures and functions in their own schema, Oracle database users need the `CREATE PROCEDURE` system privilege.

To create procedures or functions in other schemas, database users need the `CREATE ANY PROCEDURE` privilege.

To run a procedure or function, database users need the EXECUTE privilege.

Package and package body

In addition to stored procedures and functions, Oracle also provides packages to encapsulate related procedures, functions, and other program objects.

Package declares and describes all the related PL/SQL elements.

Package body contains the executable code.

To run a stored procedure or function created inside a package, specify the package name and the stored procedure or function name.

```
EXEC PKG_EMP.CALCULATE_SAL('100');
```

Examples

Create an Oracle stored procedure using the CREATE OR REPLACE PROCEDURE statement. The optional OR REPLACE clause overwrites an existing stored procedure with the same name if it exists.

```
CREATE OR REPLACE PROCEDURE EMP_SAL_RAISE
(P_EMP_ID IN NUMBER, SAL_RAISE IN NUMBER)
AS
V_EMP_CURRENT_SAL NUMBER;
BEGIN
SELECT SALARY INTO V_EMP_CURRENT_SAL FROM EMPLOYEES WHERE EMPLOYEE_ID=P_EMP_ID;
UPDATE EMPLOYEES
SET SALARY=V_EMP_CURRENT_SAL+SAL_RAISE
WHERE EMPLOYEE_ID=P_EMP_ID;
DBMS_OUTPUT.PUT_LINE('New Salary For Employee ID: '||P_EMP_ID||' Is '||(V_EMP_CURRENT_
SAL+SAL_RAISE));
EXCEPTION WHEN OTHERS THEN
RAISE_APPLICATION_ERROR(-20001,'An error was encountered - '||SQLCODE||' -ERROR-
'||SQLERRM);
ROLLBACK;
COMMIT;
END;
/
-- Run
```

```
EXEC EMP_SAL_RAISE(200, 1000);
```

Create a function using the CREATE OR REPLACE FUNCTION statement.

```
CREATE OR REPLACE FUNCTION EMP_PERIOD_OF_SERVICE_YEAR
(P_EMP_ID NUMBER)
RETURN NUMBER
AS
V_PERIOD_OF_SERVICE_YEARS NUMBER;
BEGIN
SELECT EXTRACT(YEAR FROM SYSDATE) - EXTRACT(YEAR FROM TO_DATE(HIRE_DATE)) INTO
V_PERIOD_OF_SERVICE_YEARS
FROM EMPLOYEES
WHERE EMPLOYEE_ID=P_EMP_ID;
RETURN V_PERIOD_OF_SERVICE_YEARS;
END;
/

SELECT EMPLOYEE_ID, FIRST_NAME, EMP_PERIOD_OF_SERVICE_YEAR(EMPLOYEE_ID) AS
PERIOD_OF_SERVICE_YEAR FROM EMPLOYEES;
```

EMPLOYEE_ID	FIRST_NAME	PERIOD_OF_SERVICE_YEAR
174	Ellen	13
166	Sundar	9
130	Mozhe	12
105	David	12
204	Hermann	15
116	Shelli	12
167	Amit	9
172	Elizabeth	10

Create a package using the CREATE OR REPLACE PACKAGE statement.

```
CREATE OR REPLACE PACKAGE PCK_CHINOOK_REPORTS
AS
PROCEDURE GET_ARTIST_BY_ALBUM(P_ARTIST_ID ALBUM.TITLE%TYPE);
PROCEDURE CUST_INVOICE_BY_YEAR_ANALYZE;
END;
```

Create a new package using the CREATE OR REPLACE PACKAGE BODY statement.

```
CREATE OR REPLACE PACKAGE BODY PCK_CHINOOK_REPORTS
AS
```

```
PROCEDURE GET_ARTIST_BY_ALBUM(P_ARTIST_ID ALBUM.TITLE%TYPE)
IS
V_ARTIST_NAME ARTIST.NAME%TYPE;
BEGIN
SELECT ART.NAME INTO V_ARTIST_NAME
FROM ALBUM ALB JOIN ARTIST ART USING(ARTISTID)
WHERE ALB.TITLE=P_ARTIST_ID;
DBMS_OUTPUT.PUT_LINE('ArtistName: '||V_ARTIST_NAME);
END;

PROCEDURE CUST_INVOICE_BY_YEAR_ANALYZE
AS
V_CUST_GENRES VARCHAR2(200);
BEGIN
FOR V IN(SELECT CUSTOMERID, CUSTNAME, LOW_YEAR, HIGH_YEAR, CUST_AVG FROM TMP_CUST_
INVOICE_ANALYZE)
LOOP
IF SUBSTR(V.LOW_YEAR, -4) > SUBSTR(V.HIGH_YEAR , -4) THEN
SELECT LISTAGG(GENRE, ',') WITHIN GROUP (ORDER BY GENRE) INTO V_CUST_GENRES FROM
(SELECT DISTINCT
FUNC_GENRE_BY_ID(TRC.GENREID) AS GENRE
FROM TMP_CUST_INVOICE_ANALYZE TMPTBL JOIN INVOICE INV USING(CUSTOMERID)
JOIN INVOICELINE INVLIN
ON INV.INVOICEID = INVLIN.INVOICEID
JOIN TRACK TRC
ON TRC.TRACKID = INVLIN.TRACKID
WHERE CUSTOMERID=V.CUSTOMERID);
DBMS_OUTPUT.PUT_LINE('Customer: '||UPPER(V.CUSTNAME)||' - Offer a Discount According
To Preferred Genres: '||UPPER(V_CUST_GENRES));
END IF;
END LOOP;
END;
END;

EXEC PCK_CHINOOK_REPORTS.GET_ARTIST_BY_ALBUM();
EXEC PCK_CHINOOK_REPORTS.CUST_INVOICE_BY_YEAR_ANALYZE;
```

The preceding examples demonstrate basic Oracle PL/SQL procedure and function capabilities. Oracle PL/SQL provides a large number of features and capabilities that aren't within the scope of this document.

For more information, see [CREATE FUNCTION](#) and [CREATE PROCEDURE](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL Stored Procedures provide similar functionality to Oracle stored procedures. As with Oracle, Aurora MySQL supports security execution context. It also supports input, output, and bi-directional parameters.

Stored procedures are typically used for:

- **Code reuse** — Stored procedures provide a convenient code encapsulation and reuse mechanism for multiple applications, potentially written in various languages, requiring the same database operations.
- **Security management** — By allowing access to base tables only through stored procedures, administrators can manage auditing and access permissions. This approach minimizes dependencies between application code and database code. Administrators can use stored procedures to process business rules and to perform auditing and logging.
- **Performance improvements** — Full SQL query text does not need to be transferred from the client to the database.

Note

Aurora MySQL stored procedures, triggers, and user-defined functions are collectively referred to as Stored Routines. When binary logging is enabled, MySQL SUPER privilege is required to run stored routines. However, you can run stored routines with binary logging enabled without SUPER privilege by setting the `log_bin_trust_function_creators` parameter to true for the DB parameter group for your MySQL instance.

Aurora MySQL permits stored routines to contain control flow, DML, DDL, and transaction management statements including `START TRANSACTION`, `COMMIT`, and `ROLLBACK`.

Syntax

```
CREATE [DEFINER = { user | CURRENT_USER }] PROCEDURE sp_name ([proc_parameter[,...]])
[characteristic ...]
routine_body

proc_parameter: [ IN | OUT | INOUT ] param_name type
```

```
characteristic: COMMENT 'string' | LANGUAGE SQL | [NOT] DETERMINISTIC | { CONTAINS SQL  
| NO SQL | READS SQL DATA | MODIFIES SQL DATA } | SQL SECURITY { DEFINER | INVOKER }
```

Examples

The following example demonstrates using a LOOP cursor with a source table to replace table valued parameters.

Create an OrderItems table.

```
CREATE TABLE OrderItems(  
  OrderID INT NOT NULL,  
  Item VARCHAR(20) NOT NULL,  
  Quantity SMALLINT NOT NULL,  
  PRIMARY KEY(OrderID, Item));
```

Create and populate SourceTable as a temporary data store for incoming rows.

```
CREATE TABLE SourceTable (  
  OrderID INT,  
  Item VARCHAR(20),  
  Quantity SMALLINT,  
  PRIMARY KEY (OrderID, Item));  
  
INSERT INTO SourceTable (  
  OrderID, Item, Quantity)  
VALUES (1, 'M8 Bolt', 100),  
(2, 'M8 Nut', 100),  
(3, 'M8 Washer', 200);
```

Create a procedure to loop through all rows in SourceTable and insert them into the OrderItems table.

```
CREATE PROCEDURE LoopItems()  
BEGIN  
  DECLARE done INT DEFAULT FALSE;  
  DECLARE var_OrderID INT;  
  DECLARE var_Item VARCHAR(20);  
  DECLARE var_Quantity SMALLINT;  
  DECLARE ItemCursor CURSOR FOR SELECT OrderID, Item, Quantity FROM SourceTable;  
  DECLARE CONTINUE HANDLER FOR NOT FOUND SET done = TRUE;
```

```

OPEN ItemCursor;
CursorStart: LOOP
FETCH NEXT FROM ItemCursor INTO var_OrderID, var_Item, var_Quantity;
IF Done THEN LEAVE CursorStart;
END IF;
    INSERT INTO OrderItems (OrderID, Item, Quantity)
    VALUES (var_OrderID, var_Item, var_Quantity);
END LOOP;
CLOSE ItemCursor;
END;

```

Call the stored procedure.

```
CALL LoopItems();
```

Select all rows from the OrderItems table.

```
SELECT * FROM OrderItems;
```

OrderID	Item	Quantity
1	M8 Bolt	100
2	M8 Nut	100
3	M8 Washer	200

Summary

The following table summarizes the differences between Aurora MySQL stored procedures and Oracle stored procedures.

	Oracle	Aurora MySQL	Workaround
General CREATE syntax differences	<pre>CREATE PROCEDURE <Procedure Name> Parameter1 <Type>, ...n AS <Body></pre>	<pre>CREATE PROCEDURE <Procedure Name> (Parameter1 <Type>, ...n) <Body></pre>	Rewrite stored procedure creation scripts to use PROCEDURE instead of PROC. Rewrite stored procedure creation scripts to omit the AS keyword.

	Oracle	Aurora MySQL	Workaround
Security context	<pre>{ AUTHID } { CURRENT_USER DEFINER}</pre>	<pre>DEFINER = 'user' CURRENT_USER</pre> <p>in conjunction with</p> <pre>SQL SECURITY { DEFINER INVOKER }</pre>	<p>For stored procedures that use an explicit user name, rewrite the code from EXECUTE AS 'user' to DEFINER = 'user' and SQL SECURITY DEFINER.</p> <p>For stored procedures that use the CALLER option, rewrite the code to include SQL SECURITY INVOKER.</p> <p>For stored procedures that use the SELF option, rewrite the code to DEFINER = CURRENT_USER and SQL SECURITY DEFINER.</p>
Parameter direction	IN and OUT, by default OUT can be used as IN as well.	IN, OUT, and INOUT	

For more information, see [Stored Procedures and Functions](#) and [CREATE PROCEDURE and CREATE FUNCTION Statements](#) in the *MySQL documentation*.

Oracle and MySQL regular expressions

Regular expressions help you locate and manipulate specific patterns within text data. You can leverage regular expressions for tasks such as data cleansing, validation, or transformation. The

following sections provide details on constructing and utilizing regular expressions in Oracle and MySQL with Amazon DMS.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	Syntax and option differences.

Oracle usage

A regular expression is a set of characters that define a search pattern. The most basic example is `*`, which matches any character. Most Relational Database Management Systems use the same characters for regular expressions, but some use characters differently and provide additional expressions.

Oracle SQL implementation is based on the following standards:

- IEEE Portable Operating System Interface (POSIX) standard draft 1003.2/D11.2.
- Unicode Regular Expression Guidelines of the Unicode Consortium.

Oracle SQL extends the standards as follows:

- Provides matching capabilities for multilingual data.
- Supports some commonly used PERL regular expression operators not included in the POSIX standard (for example, character class shortcuts and the non-greedy modifier `[?]`).

Summary of Oracle SQL pattern matching:

- `REGEXP_LIKE` — Can be used in `WHERE` clauses to find rows matching a certain pattern.
- `REGEXP_COUNT` — Returns the number of occurrences of a pattern in a given string.
- `REGEXP_INSTR` — Returns the position of a pattern within a string.
- `REGEXP_REPLACE` — Replaces a pattern within a string and returns the new string.

- **REGEXP_SUBSTR** — Similar to **REGEXP_INSTR**, but returns the matching substring itself instead of its position.

Summary of Oracle SQL pattern matching options:

- **i** — Case-insensitive matching.
- **c** — Case-sensitive matching.
- **n** — Allows the dot operator `.` to act like a newline character.
- **m** — Allows the string to contain multiple lines.
- **x** — Ignores white-space characters in the search pattern.

Examples

Find employees with a first name of Steven or Stephen.

```
SELECT * FROM EMPLOYEES
WHERE REGEXP_LIKE((first_name, '^Ste(v|ph)en$')
```

Find employees with a first name that includes g but not G twice starting at character position 3.

```
SELECT * FROM EMPLOYEES where
REGEXP_COUNT('George Washington', 'g', 3, 'c') = 2;
```

Find employees with a valid email address.

```
SELECT * FROM EMPLOYEES where
REGEXP_INSTR(email, '\w+@\w+(\.\w+)+') > 0;
```

Get the country with a space after each character for each employee.

```
SELECT REGEXP_REPLACE(country_name, '(.)', '\1 ') FROM EMPLOYEES;
```

For more information, see [Oracle Regular Expression Support](#) in the *Oracle documentation*.

MySQL usage

Like Oracle, Aurora MySQL Regular Expressions to make complex searches easier.

MySQL and Oracle use Henry Spencer's implementation of regular expressions, which implements the POSIX 1003.2 standard. MySQL uses the extended version to support regular expression pattern matching operations in SQL statements.

Note

Amazon Relational Database Service (Amazon RDS) for MySQL version 8.0, support for Regular Expressions will be more like Oracle. For more information, see [Regular Expressions](#) in the *MySQL documentation*.

Regular expression operators

- NOT REGEXP or NOT RLIKE — Returns 1 if the string expr does not match the regular expression specified by the pattern pat. Otherwise, it returns 0. If either expr or pat is NULL, the return value is NULL.
- REGEXP or RLIKE: Returns 1 if the string expr matches the regular expression specified by the pattern pat. Otherwise, it returns 0. If either expr or pat is NULL, the return value is NULL.

RLIKE is a synonym for REGEXP. For compatibility with Oracle, this section refers only to REGEXP.

MySQL uses the C escape syntax in strings. You must double any \ used in your REGEXP arguments.

Examples

Find employees with a first name of Steven or Stephen.

```
SELECT * FROM EMPLOYEES WHERE first_name REGEXP ('^Ste(v|ph)en$');
```

Find employees with a valid email address.

```
SELECT * FROM EMPLOYEES where  
email NOT REGEXP '^[A-Z0-9._%~]+@[A-Z0-9.-]+\.[A-Z]{2,4}$';
```

Summary

Search or usage	Oracle	MySQL
Find employees with the first name of Steven or Stephen	<pre>SELECT * FROM EMPLOYEES WHERE REGEXP_LIKE((first_name, '^Ste(v ph)en\$')</pre>	<pre>SELECT * FROM EMPLOYEES WHERE first_name REGEXP ('^Ste(v ph)en\$');</pre>
Find employees with the first name that includes g but not G twice , starting at character position 3	<pre>SELECT * FROM EMPLOYEES WHERE REGEXP_COUNT('George Washington', 'g', 3, 'c') = 2;</pre>	<pre>select * FROM EMPS WHERE LENGTH(SUBSTRING(FULL _NAME,3)) - LENGTH(REPLACE (SUBSTRING(FULL_NAM E,3), 'g', '')) = 2;</pre>
Find employees with a valid email address	<pre>SELECT * FROM EMPLOYEES where REGEXP_INSTR(email, '\w+@\w+ (\.\w+)+') >0;</pre>	<pre>SELECT * FROM EMPLOYEES where email NOT REGEXP '^[A- Z0-9._%~]+@[A-Z0-9.-]+ \.[A-Z]{2,4}\$';</pre>
Get each employee's country with space after each character	<pre>SELECT REGEXP_REPLACE (country_name, '(.)', '\1 ') FROM EMPLOYEES;</pre>	Make sure that you use a user-defined function

For more information, see [Regular Expressions](#) and [Pattern Matching](#) in the *MySQL documentation*.

Oracle TIMEZONE data type and functions and MySQL CONVERT_TZ function

With Amazon DMS, you can convert date and time values between different time zones when migrating databases. The Oracle TIMEZONE data type and functions, along with the MySQL CONVERT_TZ function, facilitate working with timestamps across time zones. The following

sections provide details on leveraging Oracle TIMEZONE and MySQL CONVERT_TZ during database migrations using Amazon DMS.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Date and Time Functions	<p>MySQL doesn't provide an equivalent option for CREATE TABLE...TIMESTAMP WITH TIMEZONE in Oracle but you can use CONVERT_TZ to achieve the same results.</p>

Oracle usage

Oracle uses data types and functions to integrate with time zones. For more information, see [A Time Zones](#) in the *Oracle documentation*.

The following data types are variants of TIMESTAMP:

- `TIMESTAMP WITH LOCAL TIMEZONE` — Data stored in the database is normalized to the database time zone, and the time zone offset is not stored as part of the column data. When users retrieve the data, Oracle returns it in the user's local session time zone.
- `TIMESTAMP WITH TIMEZONE` — Includes a time zone offset or time zone region name in its value.

Best practices:

- Use the `TIMESTAMP WITH TIMEZONE` data type when the application is used across time zones.

- The `TIMESTAMP WITH TIME ZONE` data type requires 13 bytes of storage; two more bytes of storage than `TIMESTAMP WITH LOCAL TIME ZONE` data types.

Note

The retrieved time zone offset is the difference in hours and minutes between local time and UTC.

Time zone functions

Function	Description
<code>NEW_TIME</code>	Converts date and time from one time zone to another.
<code>FROM_TZ</code>	Converts a TZ to a <code>TIMESTAMP WITH TIME ZONE</code> value.
<code>CURRENT_TIMESTAMP</code>	Returns the current date and time in the session time zone.
<code>DBTIMEZONE</code>	Returns the current date and time in the database time zone.
<code>SYS_EXTRACT_UTC</code>	Returns the UTC date and time.
<code>TO_TIMESTAMP_TZ</code>	Converts a character string of <code>CHAR</code> , <code>VARCHAR2</code> , <code>NCHAR</code> , or <code>NVARCHAR2</code> to <code>TIMESTAMP WITH TIME ZONE</code> .

Examples

Create a table using `TIMESTAMP WITH LOCAL TIME ZONE`. Note that the last inserted row is displayed as a local session timestamp. It is the only row inserted using a specific time zone that is not `LOCAL`.

```
CREATE TABLE tz_local
```

```
(id NUMBER, tz_col TIMESTAMP WITH LOCAL TIME ZONE);

INSERT INTO tz_local VALUES(1, '01-JAN-2018 2:00:00');
INSERT INTO tz_local VALUES(2, TIMESTAMP '2018-01-01 2:00:00');
INSERT INTO tz_local VALUES(3, TIMESTAMP '2018-01-01 2:00:00 -08:00');

COMMIT;

SELECT * FROM tz_local;

ID  TZ_COL
1   2018-01-01 02:00:00
2   2018-01-01 02:00:00
3   2018-01-01 05:00:00
```

Create a table using `TIMESTAMP WITH TIME ZONE`. Note that the last inserted row is displayed as a local session timestamp. It is the only row that inserted using a specific time zone.

```
ALTER SESSION SET TIME_ZONE='-4:00';
CREATE TABLE tz_tbl (id NUMBER, tz_col TIMESTAMP WITH TIME ZONE);

INSERT INTO tz_tbl VALUES(1, '01-JAN-2018 2:00:00 AM -5:00');
INSERT INTO tz_tbl VALUES(2, TIMESTAMP '2018-01-01 3:00:00');
INSERT INTO tz_tbl VALUES(3, TIMESTAMP '2018-01-01 2:00:00 -8:00');

COMMIT;

SELECT * FROM tz_tbl;

ID  TZ_COL
1   01-JAN-03 02:00:00.000000 AM -07:00
2   01-JAN-03 02:00:00.000000 AM -07:00
3   01-JAN-03 02:00:00.000000 AM -08:00
```

MySQL usage

MySQL uses time zone data type and functions similar to Oracle. Unlike Oracle, MySQL does not have many time zone options. Most functionality can be achieved when querying and not when running DDLs such as `CREATE TABLE` command in Oracle.

When the server starts, it places the host time zone in the `system_time_zone` system variable. This variable can be modified by setting the time zone operating system environment variable.

There is no equivalent option for Oracle CREATE TABLE...TIMESTAMP WITH TIME ZONE .

Comparison of time zone functions

Oracle function	MySQL function
NEW_TIME	You can use CONVERT_TZ , but you have to specify the source time zone.
FROM_TZ	CONVERT_TZ
DBTIMEZONE	CONVERT_TZ(CURRENT_TIME(),@global.time_zone,@global.time_zone)
SYS_EXTRACT_UTC	CONVERT_TZ(CURRENT_TIME(),@global.time_zone,'+00:00')
TO_TIMESTAMP_TZ	CONVERT_TZ(STR_TO_DATE('17-09-2010 23:15','%d-%m-%Y %H:%i'),@global.time_zone,'+03:00')

Examples

Query the global and session level time zone.

```
SELECT @@global.time_zone, @@session.time_zone;
```

```
@@global.time_zone  @@session.time_zone
SYSTEM              Europe/Moscow
```

For more information, see [MySQL Server Time Zone Support](#) and [Date and Time Functions](#) in the *MySQL documentation*.

Oracle and MySQL user-defined functions

The following sections will provide details on assessing, preparing, and running the migration of user-defined functions (UDFs) using Amazon DMS.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Stored Procedures	Syntax and option differences.

Oracle usage

You can create an Oracle UDF using PL/SQL, Java, or C. UDFs are useful for providing functionality not available in SQL or SQL built-in functions. They can appear in SQL statements wherever built-in SQL functions can appear.

You can use UDFs in the following cases:

- To return a single value from a SELECT statement (scalar function).
- While performing DML operations.
- In WHERE, GROUP BY, ORDER BY, HAVING, CONNECT BY, and START WITH clauses.

Examples

Create a simple Oracle UDF with arguments for employee HIRE_DATE and SALARY as INPUT parameters and calculate the overall salary over the employee's years of service for the company.

```

CREATE OR REPLACE FUNCTION TOTAL_EMP_SAL_BY_YEARS
(p_hire_date DATE, p_current_sal NUMBER)
RETURN NUMBER
AS
v_years_of_service NUMBER;
v_total_sal_by_years NUMBER;
BEGIN
SELECT EXTRACT(YEAR FROM SYSDATE) - EXTRACT(YEAR FROM to_date(p_hire_date))
INTO v_years_of_service FROM dual;
v_total_sal_by_years:=p_current_sal*v_years_of_service;
RETURN v_total_sal_by_years;
END;
/
-- Verifying

```

```
SELECT EMPLOYEE_ID, FIRST_NAME, TOTAL_EMP_SAL_BY_YEARS(HIRE_DATE, SALARY)AS TOTAL_
SALARY
FROM EMPLOYEES;
```

```
EMPLOYEE_ID FIRST_NAME TOTAL_SALARY
100          Steven      364000
101          Neena       204000
102          Lex         272000
103          Alexander  99000
104          Bruce       60000
105          David       57600
...
```

For more information, see [CREATE FUNCTION](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL supports user-defined scalar functions only. There is no support for table-valued functions.

Aurora MySQL doesn't permit stored functions to contain explicit SQL transaction statements such as COMMIT and ROLLBACK.

In Aurora MySQL, you can explicitly specify several options with the CREATE FUNCTION statement. These characteristics are saved with the function definition and are viewable with the SHOW CREATE FUNCTION statement.

- The DETERMINISTIC option must be explicitly stated. Otherwise, the engine assumes it is not deterministic.

Note

MySQL doesn't check the validity of the deterministic property declaration. If you wrongly specify a function as DETERMINISTIC when it is not, unexpected results and errors may occur.

- CONTAINS SQL indicates the function code does not contain statements that read or modify data.
- READS SQL DATA indicates the function code contains statements that read data such as SELECT but not statements that modify data such as INSERT, DELETE, or UPDATE.

- **MODIFIES SQL DATA** indicates the function code contains statements that may modify data.

Note

These options are advisory only. The server doesn't constrain the function code based on the declaration. This feature is useful for code management.

Syntax

```
CREATE FUNCTION <Function Name> ([<Function Parameter>[,...]])
RETURNS <Returned Data Type> [characteristic ...]
<Function Code Body>

characteristic:
COMMENT '<Comment>' | LANGUAGE SQL | [NOT] DETERMINISTIC
| { CONTAINS SQL | NO SQL | READS SQL DATA | MODIFIES SQL DATA }
| SQL SECURITY { DEFINER | INVOKER }
```

Migration considerations

For scalar functions, migration should be straightforward with respect to the function syntax. Rules in Aurora MySQL regarding functions are much more lenient than Oracle.

A function in Aurora MySQL may modify data and schema. Function determinism must be explicitly stated, unlike Oracle that infers it from the code. Additional properties can be stated for a function, but most are advisory only and have no functional impact.

The AS keyword, which is mandatory in Oracle before the function's code body, is not valid Aurora MySQL syntax and must be removed.

Examples

Create a scalar function to change the first character of a string to upper case.

```
CREATE FUNCTION UpperCaseFirstChar (String VARCHAR(20))
RETURNS VARCHAR(20)
BEGIN
RETURN CONCAT(UPPER(LEFT(String, 1)) , LOWER(SUBSTRING(String, 2, 19)));
END
```

```
SELECT UpperCaseFirstChar ('mIxEdCasE');
```

Mixedcase

Summary

The following table identifies similarities, differences, and key migration considerations.

Oracle	Aurora MySQL	Comment
Scalar UDF	Scalar UDF	Use CREATE FUNCTION with similar syntax, remove the AS keyword.
Inline table-valued UDF	N/A	Use views and replace parameters with WHERE filter predicates.
Multi-statement table-valued UDF	N/A	Use stored procedures to populate tables and read from the table directly.
UDF determinism implicit	Explicit declaration	Use the DETERMINISTIC characteristic explicitly to denote a deterministic function, which enables engine optimizations.
UDF boundaries local only	Can change data and schema	UDF rules are more lenient, avoid unexpected changes from function invocation.

For more information, see [CREATE PROCEDURE and CREATE FUNCTION Statements](#) in the *MySQL documentation*.

Oracle UTL_FILE and MySQL integration with Amazon S3

With Amazon DMS, you can seamlessly migrate Oracle databases utilizing UTL_FILE and MySQL databases with Amazon S3 integration to Amazon. The following sections outline the steps to configure and utilize UTL_FILE with Oracle and MySQL integration with Amazon S3 through Amazon DMS.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	MySQL doesn't support UTL_FILE but Aurora MySQL has a built-in integration with Amazon S3.

Oracle usage

Oracle UTL_FILE PL/SQL package enables you to access files stored outside of the database such as files stored on the operating system, the database server, or a connected storage volume. UTL_FILE.FOPEN, UTL_FILE.GET_LINE, and UTL_FILE.PUT_LINE are procedures within the UTL_FILE package used to open, read, and write files.

Examples

Run an anonymous PL/SQL block that reads a single line from file1 and writes it to file2.

- Use UTL_FILE.FILE_TYPE to create a handle for the file.
- Use UTL_FILE.FOPEN to open stream access to the file and specify:
 - The logical Oracle directory object pointing to the O/S folder where the file resides.
 - The file name.
 - The file access mode: 'A'=append mode, 'W'=write mode
- Use UTL_FILE.GET_LINE to read a line from the input file into a variable.

- Use `UTL_FILE.PUT_LINE` to write a single line to the output file.

```
DECLARE
strString1 VARCHAR2(32767);
fileFile1 UTL_FILE.FILE_TYPE;
BEGIN
fileFile1 := UTL_FILE.FOPEN('FILES_DIR', 'File1.tmp', 'R');
UTL_FILE.GET_LINE(fileFile1, strString1);
UTL_FILE.FCLOSE(fileFile1);
fileFile1 := UTL_FILE.FOPEN('FILES_DIR', 'File2.tmp', 'A');
utl_file.PUT_LINE(fileFile1, strString1);
utl_file.fclose(fileFile1);
END;
/
```

For more information, see [UTL_FILE](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL provides similar functionality to Oracle `UTL_FILE` with Amazon S3 integration.

There two important integration aspects between Aurora MySQL and Amazon S3:

- Saving data to an S3 file.
- Loading data from an S3 file.

Note

Make sure that Aurora MySQL has permissions to the S3 bucket.

Saving data to Amazon S3

You can use the `SELECT INTO OUTFILE S3` statement to query data from an Amazon Aurora MySQL DB cluster and save it directly to text files stored in an Amazon S3 bucket. You can use this approach to avoid transferring data first to the client and then copying the data from the client to Amazon S3.

Note

The default file size threshold is 6 GB. If the data selected by the statement is less than the file size threshold, a single file is created. Otherwise, multiple files are created.

If the SELECT statement failed, files already uploaded to Amazon S3 remain in the specified Amazon S3 bucket. You can use another statement to upload the remaining data instead of starting over.

If the amount of data to be selected is more than 25 GB, it is recommended to use multiple SELECT INTO OUTFILE S3 statements to save the data to Amazon S3.

Metadata, such as table schema or file metadata, isn't uploaded by Aurora MySQL to Amazon S3.

Examples

The following statement selects all data in the employees table and saves it to an Amazon S3 bucket in a different region from the Aurora MySQL DB cluster. The statement creates data files in which each field is terminated by a comma , character and each row is terminated by a newline \n character. The statement returns an error if files that match the sample_employee_data file prefix already exist in the specified Amazon S3 bucket.

```
SELECT * FROM employees INTO OUTFILE S3
's3-us-west-2://aurora-select-into-s3-pdx/sample_employee_data'
FIELDS TERMINATED BY ',' LINES TERMINATED BY '\n';
```

The following statement selects all data in the employees table and saves the data to an Amazon S3 bucket in the same region as the Aurora MySQL DB cluster. The statement creates data files in which each field is terminated by a comma , character and each row is terminated by a newline \n character. It also creates a manifest file. The statement returns an error if files that match the sample_employee_data file prefix already exist in the specified Amazon S3 bucket.

```
SELECT * FROM employees INTO OUTFILE S3
's3://aurora-select-into-s3-pdx/sample_employee_data'
FIELDS TERMINATED BY ',' LINES TERMINATED BY '\n'
MANIFEST ON;
```

The following statement selects all data in the employees table and saves the data to an Amazon S3 bucket in a different region from the Aurora database cluster. The statement creates data files

in which each field is terminated by a comma , character and each row is terminated by a newline \n character. The statement overwrites any existing files that match the sample_employee_data file prefix in the specified Amazon S3 bucket.

```
SELECT * FROM employees INTO OUTFILE S3 '  
s3-us-west-2://aurora-select-into-s3-pdx/sample_employee_data'  
FIELDS TERMINATED BY ',' LINES TERMINATED BY '\n' OVERWRITE ON;
```

The following statement selects all data in the employees table and saves the data to an Amazon S3 bucket in the same region as the Aurora MySQL DB cluster. The statement creates data files in which each field is terminated by a comma , character and each row is terminated by a newline \n character. It also creates a manifest file. The statement overwrites any existing files that match the sample_employee_data file prefix in the specified Amazon S3 bucket.

```
SELECT * FROM employees INTO OUTFILE S3  
's3://aurora-select-into-s3-pdx/sample_employee_data'  
FIELDS TERMINATED BY ',' LINES TERMINATED BY '\n'  
MANIFEST ON OVERWRITE ON;
```

For more information, see [Saving data from an Amazon Aurora MySQL DB cluster into text files in an Amazon S3 bucket](#) in the *User Guide for Aurora*.

Load XML from Amazon S3

Use the LOAD DATA FROM S3 or LOAD XML FROM S3 statement to load data from files stored in an Amazon S3 bucket.

The LOAD DATA FROM S3 statement can load data from any text file format supported by the MySQL LOAD DATA INFILE statement such as comma-delimited text data. Compressed files are not supported.

Examples

The following example runs the LOAD DATA FROM S3 statement with the manifest file named customer.manifest. After the statement completes, an entry for each successfully loaded file is written to the aurora_s3_load_history table.

```
LOAD DATA FROM S3 MANIFEST  
's3-us-west-2://aurora-bucket/customer.manifest'
```

```
INTO TABLE CUSTOMER FIELDS TERMINATED BY ','
LINES TERMINATED BY '\n'
(ID, FIRSTNAME, LASTNAME, EMAIL);
```

Every successful `LOAD DATA FROM S3` statement updates the `aurora_s3_load_history` table in the `mysql` schema with an entry for each file that was loaded.

After you run the `LOAD DATA FROM S3` statement, you can verify which files were loaded by querying the `aurora_s3_load_history` table. To see the files that were loaded from one execution of the statement, use the `WHERE` clause to filter the records on the Amazon S3 URI for the manifest file used in the statement. If you have used the same manifest file before, filter the results using the timestamp field.

```
select * from mysql.aurora_s3_load_history where load_prefix = 'S3_URI';
```

The following table describes the fields in the `aurora_s3_load_history` table.

Field	Description
<code>load_prefix</code>	<p>The URI specified in the load statement. This URI can map to any of the following:</p> <ul style="list-style-type: none"> A single data file for a <code>LOAD DATA FROM S3 FILE</code> statement. An Amazon S3 prefix that maps to multiple data files for a <code>LOAD DATA FROM S3 PREFIX</code> statement. A single manifest file containing the names of files to be loaded for a <code>LOAD DATA FROM S3 MANIFEST</code> statement.
<code>file_name</code>	<p>The name of a file that was loaded into Aurora from Amazon S3 using the URI identified in the <code>load_prefix</code> field.</p>
<code>version_number</code>	<p>The version number of the file identified by the <code>file_name</code> field that was loaded if the Amazon S3 bucket has a version number.</p>

Field	Description
bytes_loaded	The size in bytes of the file loaded.
load_timestamp	The timestamp when the LOAD DATA FROM S3 statement completed.

The following statement loads data from an Amazon S3 bucket in the same region as the Aurora DB cluster. It reads the comma-delimited data in the `customerdata.txt` file residing in the `dbbucket` Amazon S3 bucket and then loads the data into the table `store-schema.customer-table`.

```
LOAD DATA FROM S3 's3://dbbucket/customerdata.csv'
INTO TABLE store-schema.customer-table
FIELDS TERMINATED BY ',' LINES TERMINATED BY '\n'
(ID, FIRSTNAME, LASTNAME, ADDRESS, EMAIL, PHONE);
```

The following statement loads data from an Amazon S3 bucket in a different region from the Aurora DB cluster. The statement reads the comma-delimited data from all files that match the `employee-data` object prefix in the `my-data` Amazon S3 bucket in the `us-west-2` region and then loads the data into the `employees` table.

```
LOAD DATA FROM S3 PREFIX
's3-us-west-2://my-data/employee_data'
INTO TABLE employees
FIELDS TERMINATED BY ','
LINES TERMINATED BY '\n'
(ID, FIRSTNAME, LASTNAME, EMAIL, SALARY);
```

The following statement loads data from the files specified in a JSON manifest file named `q1_sales.json` into the `sales` table.

```
LOAD DATA FROM S3 MANIFEST
's3-us-west-2://aurora-bucket/q1_sales.json'
INTO TABLE sales FIELDS TERMINATED BY ','
LINES TERMINATED BY '\n' (MONTH, STORE, GROSS, NET);
```

You can use the `LOAD XML FROM S3` statement to load data from XML files stored on an Amazon S3 bucket in one of three different XML formats as described below.

Column names as attributes of a `<row>` element. The attribute value identifies the contents of the table field.

```
<row column1="value1" column2="value2" .../>
```

Column names as child elements of a `<row>` element. The value of the child element identifies the contents of the table field.

```
<row>
<column1>value1</column1>
<column2>value2</column2>
</row>
```

Column names in the name attribute of `<field>` elements in a `<row>` element. The value of the `<field>` element identifies the contents of the table field.

```
<row>
<field name='column1'>value1</field>
<field name='column2'>value2</field>
</row>
```

The following statement loads the first column from the input file into the first column of table1 and sets the value of the table_column2 column in table1 to the input value of the second column divided by 100.

```
LOAD XML FROM S3 's3://mybucket/data.xml'
INTO TABLE table1 (column1, @var1)
SET table_column2 = @var1/100;
```

The following statement sets the first two columns of table1 to the values in the first two columns from the input file and then sets the value of the column3 in table1 to the current time stamp.

```
LOAD XML FROM S3 's3://mybucket/data.xml'
INTO TABLE table1 (column1, column2)
SET column3 = CURRENT_TIMESTAMP;
```

You can use subqueries in the right side of SET assignments. For a subquery that returns a value to be assigned to a column, you can use only a scalar subquery. Also, you cannot use a subquery to select from the table that is being loaded.

For more information, see [Loading data into an Amazon Aurora MySQL DB cluster from text files in an Amazon S3 bucket](#) in the *User Guide for Aurora*.

Oracle UTL_MAIL or UTL_SMTP and Amazon Simple Notification Service

With Amazon DMS, you can migrate email functionality from Oracle databases to Amazon Simple Notification Service (Amazon SNS). Oracle UTL_MAIL and UTL_SMTP packages provide database email capabilities, which Amazon Database Migration Service can help you transition to the fully managed Amazon SNS service. The following sections detail the process of replicating Oracle UTL_MAIL/UTL_SMTP functionality using Amazon SNS with Amazon DMS.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	Use Lambda integration.

Oracle UTL_MAIL usage

The Oracle UTL_MAIL package provides functionality for sending email messages. Unlike UTL_SMTP, which is more complex and provided in earlier versions of Oracle, UTL_MAIL supports attachments. For most cases, UTL_MAIL is a better choice.

Examples

Install the required mail packages.

```
@{ORACLE_HOME}/rdbms/admin/utlmail.sql
@{ORACLE_HOME}/rdbms/admin/prvtmail.plb
```

Set the smtp_out_server parameter.

```
ALTER SYSTEM SET smtp_out_server = 'smtp.domain.com' SCOPE=BOTH;
```

Send an email message.

```
exec utl_mail.send('Sender@mailserver.com', 'recipient@mailserver.com', NULL, NULL,
  'This is the subject', 'This is the message body', NULL, 3, NULL);
```

For more information, see [UTL_MAIL](#) in the *Oracle documentation*.

Oracle UTL_SMTP usage

The Oracle UTL_SMTP package provides functionality for sending email messages and is useful for sending alerts about database events. Unlike UTL_MAIL, UTL SMTP is more complex and doesn't support attachments. For most cases, UTL_MAIL is a better choice.

Examples

The following example demonstrates using UTL_SMTP procedures to send email messages.

Install the required scripts.

```
In oracle 12c:
@{ORACLE_HOME}/rdbms/admin/utlsmtp.sql

In oracle 11g:
@{ORACLE_HOME}/javavm/install/initjvm.sql
@{ORACLE_HOME}/rdbms/admin/initplsjsql
```

Create and send an email message.

- UTL_SMTP.OPEN_CONNECTION opens a connection to the smtp server.
- UTL_SMTP.HELO initiates a handshake with the smtp server.
- UTL_SMTP.MAIL Initiates a mail transaction that obtains the senders details.
- UTL_SMTP.RCPT adds a recipient to the mail transaction.
- UTL_SMTP.DATA adds the message content.
- UTL_SMTP.QUIT terminates the SMTP transaction.

```
DECLARE
smtpconn utl_smtp.connection;
```

```
BEGIN
smtpconn := UTL_SMTP.OPEN_CONNECTION('smtp.mailserver.com', 25);
UTL_SMTP.HELO(smtpconn, 'smtp.mailserver.com');
UTL_SMTP.MAIL(smtpconn, 'sender@mailserver.com');
UTL_SMTP.RCPT(smtpconn, 'recipient@mailserver.com');
UTL_SMTP.DATA(smtpconn, 'Message body');
UTL_SMTP.QUIT(smtpconn);
END;
/
```

For more information, see [Managing Resources with Oracle Database Resource Manager](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL does not support direct configuration of engine alerts. Use the Event Notifications Infrastructure to collect history logs or receive event notifications in near real-time.

The Amazon Relational Database Service (Amazon RDS) uses the Amazon Simple Notification Service (Amazon SNS) to provide notifications for events. Amazon SNS can send notifications in any form supported by the region including email, text messages, or calls to HTTP endpoints for response automation.

Events are grouped into categories. You can only subscribe to event categories, not individual events. SNS sends notifications when any event in a category occurs.

You can subscribe to alerts for database instances, database clusters, database snapshots, database cluster snapshots, database security groups, and database parameter groups. For example, a subscription to the Backup category for a specific database instance sends notifications when backup-related events occur on that instance. A subscription to the Configuration Change category for a database security group sends notifications when the security group changes.

Note

For Amazon Aurora, some events occur at the cluster rather than instance level. You will not receive those events if you subscribe to an Aurora DB instance.

Amazon SNS sends event notifications to the address specified when the subscription was created. Typically, administrators create several subscriptions. For example, one subscription to receive

logging events and another to receive only critical events for a production environment requiring immediate responses.

You can disable notifications without deleting a subscription by setting the Enabled radio button to No in the Amazon RDS console. Alternatively, use the Amazon Command Line Interface (CLI) or Amazon RDS API to change the Enabled setting.

Subscriptions are identified by the Amazon Resource Name (ARN) of an Amazon SNS topic. The Amazon RDS console creates ARNs when subscriptions are created. When using the CLI or API, you must create the ARN using the Amazon SNS console or the Amazon SNS API.

Examples

The following walkthrough demonstrates how to create an Event Notification Subscription:

1. Sign in to the Amazon Management Console and choose **RDS**.
2. Choose **Events**. If you have not previously subscribed to events, the screen displays zero events.
3. Choose **Event subscriptions**, and then choose **Create event subscription**.
4. For **Name**, enter the name of the subscription.
5. For **Target**, choose **ARN** or **New email topic**. For email subscriptions, enter values for **Topic name** and **With these recipients**.
6. Choose the event source and then choose specific event categories to be monitored from the drop-down menu.
7. Choose **Create**.
8. On the Amazon RDS dashboard, choose **Recent events**.

For more information, see [Working with Amazon RDS event notification](#) in the *Amazon Relational Database Service User Guide*.

For application email requirements, consider using a dedicated email framework. If the code generating email messages must reside in the database, consider using a queue table. Replace all occurrences of UTL_SMTP and UTL_MAIL with an INSERT into the queue table. Design external applications to connect, read the queue, send an email message, and then update the status periodically. With this approach, messages can be populated with a query result similar to UTL_SMTP and UTL_MAIL with the query option.

The only way to send email from the database is to use Amazon Lambda integration. For more information about Amazon Lambda, see [Amazon Lambda](#).

For an example of sending an email message from Aurora MySQL using Amazon Lambda integration, see [Invoking a Lambda Function from an Amazon Aurora MySQL DB Cluster](#).

Oracle and MySQL tables and indexes

This section provides reference pages for Oracle and MySQL tables and indexes.

Topics

- [Case sensitivity differences for Oracle and MySQL](#)
- [Data types](#)
- [Oracle Read-only tables and partitions and Amazon Aurora MySQL replicas](#)
- [Oracle and MySQL table constraints](#)
- [Oracle and MySQL temporary tables](#)
- [Oracle and MySQL triggers](#)
- [Oracle and MySQL tablespaces and data files](#)
- [Oracle user-defined types](#)
- [Oracle unused columns](#)
- [Oracle virtual columns and MySQL generated columns](#)
- [MySQL overall indexes summary](#)
- [Oracle bitmap indexes](#)
- [Oracle and MySQL B-tree indexes](#)
- [Oracle composite indexes and MySQL multiple-column indexes](#)
- [Oracle function-based indexes and MySQL indexing on generated columns](#)
- [Oracle and MySQL invisible indexes](#)
- [Oracle index-organized table and MySQL InnoDB clustered index](#)
- [Oracle local and global partitioned indexes and MySQL partitioned indexes](#)
- [Oracle automatic indexing](#)

Case sensitivity differences for Oracle and MySQL

Object name case sensitivity is different for Oracle and MySQL. Oracle names aren't case sensitive. Aurora MySQL names are case sensitive.

In Aurora for MySQL, the case sensitivity is determined by the value of the `lower_case_table_names` parameter. You can choose one of the three possible values for this parameter. To avoid some issues, Amazon recommends to use only two values with this parameter:

- 0 (names stored as given and comparisons are case-sensitive) is supported for all Amazon RDS for MySQL versions.
- 1 (names stored in lowercase and comparisons are not case-sensitive) is supported for Amazon RDS for MySQL version 5.6, version 5.7, and version 8.0.19 and higher 8.0 versions.

The `lower_case_table_names` parameter should be set as part of a custom DB parameter group before creating a DB instance. You should avoid changing the `lower_case_table_names` parameter for existing database instances because doing so could cause inconsistencies with point-in-time recovery backups and read replica DB instances.

Read replicas should always use the same `lower_case_table_names` parameter value as the source DB instance.

By default, object names are being stored in lowercase for MySQL. In most cases, you'll want to use Amazon Database Migration Service (Amazon DMS) transformations to change schema, table, and column names to lowercase.

For example, to create a table named `EMPLOYEES` (uppercase) in MySQL, you should use double quotation marks as shown in the following code example. You can use the same approach in Oracle.

```
CREATE TABLE "EMPLOYEES" (  
  EMP_ID NUMERIC PRIMARY KEY,  
  EMP_FULL_NAME VARCHAR(60) NOT NULL,  
  AVG_SALARY NUMERIC NOT NULL);
```

The following command creates a table named `employees` in lowercase.

```
CREATE TABLE EMPLOYEES (  
  EMP_ID NUMERIC PRIMARY KEY,  
  EMP_FULL_NAME VARCHAR(60) NOT NULL,  
  AVG_SALARY NUMERIC NOT NULL);
```

MySQL will look for objects names in with the exact case sensitivity as written in the query.

You can disable table name case sensitivity in MySQL by setting the parameter `lower_case_table_names` to 1. Column, index, stored routine, event names, and column aliases are not case sensitive on either platform.

For more information, see [Identifier Case Sensitivity](#) in the *MySQL documentation*.

Data types

With Amazon DMS, you can migrate data between different database platforms, allowing you to consolidate databases, perform database modernization, or migrate databases to the cloud. Data types define the kind of data that can be stored in a database column or variable. The following sections will provide detailed information about different data types supported by Oracle and MySQL.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Data Types	Aurora MySQL doesn't support BFILE, ROWID, and UROWID.

Oracle usage

Oracle provides a set of primitive data types for defining table columns and PL/SQL code variables. The assigned data types for table columns or PL/SQL code (such as stored procedures and triggers) define the valid values each column or argument can store.

Oracle data types and MySQL data types

Character data types

Oracle data type	Oracle data type characteristic	MySQL identical compatibility	MySQL corresponding data type
CHAR(n)	Maximum size of 2000 bytes	Yes	CHAR(n)
CHARACTER(n)	Maximum size of 2000 bytes	Yes	CHARACTER(n)
NCHAR(n)	Maximum size of 2000 bytes	Yes	NCHAR(n)
VARCHAR(n)	Maximum size of 2000 bytes	Yes	VARCHAR(n)
NCHAR VARYING(n)	Varying-length UTF-8 string, maximum size of 4000 bytes	Yes	NCHAR VARYING(n)
VARCHAR2(n) 11g	Maximum size of 4000 bytes or 32 KB in PL/SQL	No	VARCHAR(n)
VARCHAR2(n) 12g	Maximum size of 32767 bytes MAX_STRING_SIZE=EXTENDED	No	VARCHAR(n)
NVARCHAR2(n)	Maximum size of 4000 bytes	No	VARCHAR(n)
LONG	Maximum size of 2 GB	Yes	LONG
RAW(n)	Maximum size of 2000 bytes	No	VARBINARY(n)

Oracle data type	Oracle data type characteristic	MySQL identical compatibility	MySQL corresponding data type
LONG RAW	Maximum size of 2 GB	No	LONGTEXT

Numeric data types

Oracle data type	Oracle data type characteristic	MySQL identical compatibility	MySQL corresponding data type
NUMBER	Floating-point number	No	DECIMAL(p, s)
NUMBER(*)	Floating-point number	No	DOUBLE
NUMBER(p, s)	Precision can range from 1 to 38, scale can range from -84 to 127	No	DECIMAL(p, s)
NUMERIC(p, s)	Precision can range from 1 to 38	Yes	NUMERIC(p, s)
FLOAT(p)	Floating-point number	Yes	FLOAT(p)
DEC(p, s)	Fixed-point number	Yes	DEC(p, s)
DECIMAL(p, s)	Fixed-point number	Yes	DECIMAL(p, s)
INT	38 digits integer	Yes	INT
INTEGER	38 digits integer	Yes	INTEGER
SMALLINT	38 digits integer	Yes	SMALLINT

Oracle data type	Oracle data type characteristic	MySQL identical compatibility	MySQL corresponding data type
REAL	Floating-point number	Yes	REAL
DOUBLE PRECISION	Floating-point number	Yes	DOUBLE PRECISION

Date and time data types

Oracle data type	Oracle data type characteristic	MySQL identical compatibility	MySQL corresponding data type
DATE	Stores date and time data (year, month, day, hour, minute and second)	Yes	DATETIME
TIMESTAMP(p)	Date and time with fraction	Yes	TIMESTAMP(6)
TIMESTAMP(p) WITH TIME ZONE	Date and time with fraction and time zone	No	DATETIME(n)
INTERVAL YEAR(p) TO MONTH	Date interval	No	VARCHAR(n)
INTERVAL DAY(p) TO SECOND(s)	Day and time interval	No	VARCHAR(n)

LOB data types

Oracle data type	Oracle data type characteristic	MySQL identical compatibility	MySQL corresponding data type
BFILE	Pointer to binary file, maximum file size of 4 GB	No	VARCHAR (255)
BLOB	Binary large object, maximum file size of 4 GB	Yes	BLOB
CLOB	Character large object, maximum file size of 4 GB	No	LONGTEXT
NCLOB	Variable-length Unicode string, maximum file size of 4 GB	No	LONGTEXT

ROWID data types

Oracle data type	Oracle data type characteristic	MySQL identical compatibility	MySQL corresponding data type
ROWID	Physical row address	No	CHAR(n)
UROWID(n)	Universal row id, logical row addresses	No	VARCHAR(n)

XML data type

Oracle data type	Oracle data type characteristic	MySQL identical compatibility	MySQL corresponding data type
XMLTYPE	XML data	No	LONGTEXT

Logical data type

Oracle data type	Oracle data type characteristic	MySQL identical compatibility	MySQL corresponding data type
BOOLEAN	Values TRUE, FALSE, and NULL, can't be assigned to a database table column	Yes	BOOLEAN

Spatial types

Oracle data type	Oracle data type characteristic	MySQL identical compatibility	MySQL corresponding data type
SDO_GEOMETRY	The geometric description of a spatial object	No	N/A
SDO_TOPO_GEOMETRY	Describes a topology geometry	No	N/A
SDO_GEORASTER	A raster grid or image object is stored in a single row	No	N/A

Media types

Oracle data type	Oracle data type characteristic	MySQL identical compatibility	MySQL corresponding data type
ORDDiCom	Supports the storage and management of audio data	No	N/A

Oracle data type	Oracle data type characteristic	MySQL identical compatibility	MySQL corresponding data type
ORDDicom	Supports the storage and management of Digital Imaging and Communications in Medicine (DICOM).	No	N/A
ORDDoc	Supports storage and management of any type of media data	No	N/A
ORDImage	Supports the storage and management of image data	No	N/A
ORDVideo	Supports the storage and management of video data	No	N/A

Oracle character column semantics

Oracle supports BYTE and CHAR semantics for column size, which determines the amount of storage allocated for CHAR and VARCHAR columns.

- If you define a field as `VARCHAR2(10 BYTE)`, Oracle can use up to 10 bytes for storage. However, based on your database codepage and NLS settings, you may not be able to store 10 characters in that field because the physical size of some non-English characters exceeds one byte.
- If you define a field as `VARCHAR2(10 CHAR)`, Oracle can store 10 characters no matter how many bytes are required to store each non-English character.

```
CREATE TABLE table1 (col1 VARCHAR2(10 CHAR), col2 VARCHAR2(10 BYTE));
```

By default, Oracle uses BYTE semantics. When using a multi-byte character set such as UTF8, use one of the following options.

- Use the CHAR modifier in the VARCHAR2 or CHAR column definition.
- Modify the session or system parameter NLS_LENGTH_SEMANTICS to change the default from BYTE to CHAR.

```
ALTER system SET nls_length_semantics=char scope=both;
ALTER system SET nls_length_semantics=byte scope=both;

ALTER session SET nls_length_semantics=char;
ALTER session SET nls_length_semantics=byte;
```

For more information, see [Data Types](#) in the *Oracle documentation*.

MySQL usage

MySQL provides multiple data types equivalent to certain Oracle data types. The following table provides the full list of MySQL data types.

Character data types

MySQL data type	MySQL data type characteristic
CHAR(n)	Stores exactly (n) characters.
VARCHAR(n)	Stores a variable number of characters, up to a maximum of n characters.
BINARY	Stores exactly (n) bytes.
VARBINARY	Stores a variable number of characters, up to a maximum of n bytes.
BLOB	Binary large object that can hold a variable amount of data.

MySQL data type	MySQL data type characteristic
TEXT	Specific variant of varchar, which does not require you to specify an upper limit on the number of characters.
ENUM	String object with a value chosen from a list of permitted values that are enumerated explicitly in the column specification at table creation time.
SET	String object that can have zero or more values, each of which must be chosen from a list of permitted values specified when the table is created.

Numeric data types

MySQL data type	MySQL data type characteristic
INTEGER	Max value is 2147483647.
INT	Max value is 2147483647.
SMALLINT	Max value is 32767.
TINYINT	Max value is 127.
MEDIUMINT	Max value is 8388607.
BIGINT	Max value is $2^{63}-1$.
DECIMAL (p, s)	Stores any value with p digits and s decimals.
NUMERIC(p, s)	Stores any value with p digits and s decimals.
FLOAT (m, d)	Values can be stored with up to M digits in total, of which D digits may be after the decimal point.

MySQL data type	MySQL data type characteristic
DOUBLE (m, d)	Values can be stored with up to M digits in total, of which D digits may be after the decimal point.
BIT (m)	Stores M-bit values. M can range from 1 to 64.

Date and time data types

MySQL data type	MySQL data type characteristic
DATE	Values with a date part but no time part. MySQL retrieves and displays DATE values in the YYYY-MM-DD format. The supported range is 1000-01-01 to 9999-12-31 .
DATETIME	Values that contain both date and time parts. MySQL retrieves and displays DATETIME values in YYYY-MM-DD HH:MM:SS format. The supported range is 1000-01-01 00:00:00 to 9999-12-31 23:59:59 .
TIMESTAMP	Values that contain both date and time parts. TIMESTAMP has a range of 1970-01-01 00:00:01 UTC to 2038-01-19 03:14:07 UTC.
TIME	Values may range from -838:59:59 to 838:59:59 . The hours part may be so large because the TIME type can be used not only to represent a time of day, which must be less than 24 hours, but also elapsed time or a time interval between two events, which may be much greater than 24 hours, or even negative.

MySQL data type	MySQL data type characteristic
YEAR	YEAR 1-byte type used to represent year values. It can be declared as YEAR or YEAR(n) and has a display width of n characters.

Logical data type

MySQL data type	MySQL data type characteristic
BOOLEAN	Holds a truth value. Will accept values such as TRUE, 't','true', 'y', 'yes', and '1' as true. Uses 1 byte of storage, and can store NULL.

Geometric data types

MySQL data type	MySQL data type characteristic
GEOMETRY	The column type to specify when you want to use the data models below.
POINT	An (x,y) value.
LINestring	A line (pt1, pt2).
POLYGON	A sequence of points, effectively a closed path.
MULTIPOINT	Collection of POINTs.
MULTILINestring	Collection of LINEs.
MULTIPOLYGON	Collection of POLYGONs.
GEOMETRYCOLLECTION	Collection of geometry data types.

Other data types

MySQL data type	MySQL data type characteristic
JSON	Textual JSON data

Migration of Oracle data types to MySQL data types

You can perform automatic migration and conversion of Oracle tables and data types using Amazon Schema Conversion Tool (Amazon SCT).

Examples

To demonstrate Amazon SCT capability for migrating Oracle tables to their MySQL equivalents, a table containing columns representing the majority of Oracle data types was created and converted using Amazon SCT.

Source Oracle compatible DDL for creating the DATATYPES table.

```
CREATE TABLE "DATATYPES"(
  "BFILE"                BFILE,
  "BINARY_FLOAT"        BINARY_FLOAT,
  "BINARY_DOUBLE"       BINARY_DOUBLE,
  "BLOB"                 BLOB,
  "CHAR"                CHAR(10 BYTE),
  "CHARACTER"           CHAR(10 BYTE),
  "CLOB"                CLOB,
  "NCLOB"               NCLOB,
  "DATE"                DATE,
  "DECIMAL"             NUMBER(3,2),
  "DEC"                 NUMBER(3,2),
  "DOUBLE_PRECISION"    FLOAT(126),
  "FLOAT"               FLOAT(3),
  "INTEGER"             NUMBER(*,0),
  "INT"                 NUMBER(*,0),
  "INTERVAL_YEAR"       INTERVAL YEAR(4) TO MONTH,
  "INTERVAL_DAY"        INTERVAL DAY(4) TO SECOND(4),
  "LONG"                LONG,
  "NCHAR"               NCHAR(10),
  "NCHAR_VARYING"       NVARCHAR2(10),
  "NUMBER"              NUMBER(9,9),
  "NUMBER1"             NUMBER(9,0),
  "NUMBER(*)"           NUMBER,
```

```

"NUMERIC"          NUMBER(9,9),
"NVARCHAR2"       NVARCHAR2(10),
"RAW"             RAW(10),
"REAL"           FLOAT(63),
"ROW_ID"         ROWID,
"SMALLINT"       NUMBER(*,0),
"TIMESTAMP"      TIMESTAMP(5),
"TIMESTAMP_WITH_TIME_ZONE"  TIMESTAMP(5) WITH TIME ZONE,
"UROWID"         UROWID(10),
"VARCHAR"        VARCHAR2(10 BYTE),
"VARCHAR2"       VARCHAR2(10 BYTE),
"XMLTYPE"        XMLTYPE
);

```

Target MySQL compatible DDL for creating the DATATYPES table migrated from Oracle with Amazon SCT.

```

CREATE TABLE IF NOT EXISTS datatypes(
  bfile VARCHAR(1000) DEFAULT NULL,
  BINARY_FLOAT FLOAT(12) DEFAULT NULL,
  BINARY_DOUBLE DOUBLE DEFAULT NULL,
  `BLOB` LONGBLOB DEFAULT NULL,
  `CHAR` CHAR(10) DEFAULT NULL,
  `CHARACTER` CHAR(10) DEFAULT NULL,
  CLOB LONGTEXT DEFAULT NULL,
  NCLOB LONGTEXT DEFAULT NULL,
  `DATE` DATETIME DEFAULT NULL,
  `DECIMAL` DECIMAL(3,2) DEFAULT NULL,
  `DEC` DECIMAL(3,2) DEFAULT NULL,
  DOUBLE_PRECISION DOUBLE DEFAULT NULL,
  `FLOAT` DOUBLE DEFAULT NULL,
  `INTEGER` DECIMAL(38,0) DEFAULT NULL,
  `INT` DECIMAL(38,0) DEFAULT NULL,
  INTERVAL_YEAR VARCHAR(30) DEFAULT NULL,
  INTERVAL_DAY VARCHAR(30) DEFAULT NULL,
  `LONG` LONGTEXT DEFAULT NULL,
  NCHAR CHAR(10) DEFAULT NULL,
  NCHAR_VARYING VARCHAR(10) DEFAULT NULL,
  NUMBER DECIMAL(9,9) DEFAULT NULL,
  NUMBER1 DECIMAL(9,0) DEFAULT NULL,
  `NUMBER(*)` DOUBLE DEFAULT NULL,
  `NUMERIC` DECIMAL(9,9) DEFAULT NULL,
  NVARCHAR2 VARCHAR(10) DEFAULT NULL,

```

```
RAW VARBINARY(10) DEFAULT NULL,  
`REAL` DOUBLE DEFAULT NULL,  
ROW_ID CHAR(10) DEFAULT NULL,  
`SMALLINT` DECIMAL(38,0) DEFAULT NULL,  
`TIMESTAMP` DATETIME(5) DEFAULT NULL,  
TIMESTAMP_WITH_TIME_ZONE DATETIME(5) DEFAULT NULL,  
UROWID VARCHAR(10) DEFAULT NULL,  
`VARCHAR` VARCHAR(10) DEFAULT NULL,  
VARCHAR2 VARCHAR(10) DEFAULT NULL,  
XMLTYPE LONGTEXT DEFAULT NULL);
```

Amazon SCT converted most of the data types. However, a few exceptions were raised for data types that Amazon SCT is unable to automatically convert and where Amazon SCT recommended manual actions.

MySQL doesn't have a data type BFILE

BFILEs are pointers to binary files.

Recommended actions: Either store a named file with the data and create a routine that gets that file from the file system, or store the data blob inside your database.

MySQL doesn't have a data type ROWID

ROWIDs are physical row addresses inside Oracle storage subsystems. The ROWID data type is primarily used for values returned by the ROWID pseudocolumn.

Recommended actions: Although MySQL contains a `ctid` column that is the physical location of the row version within its table, it doesn't have a comparable data type. However, you can use CHAR as a partial data type equivalent. If you use ROWID data types in your code, modifications may be necessary.

MySQL doesn't have a data type UROWID

Universal row identifier, or UROWID, is a single Oracle data type that supports both logical and physical row identifiers of foreign table row identifiers such as non-Oracle tables accessed through a gateway.

Recommended actions: MySQL doesn't have a comparable data type. You can use VARCHAR(n) as a partial data type equivalent. However, if you are using UROWID data types in your code, modifications may be necessary.

For more information, see [Schema Conversion Tool Documentation](#) and [Data Types](#) in the *MySQL documentation*.

Oracle Read-only tables and partitions and Amazon Aurora MySQL replicas

With Amazon DMS, you can migrate data from Oracle databases to Aurora MySQL databases while maintaining read-only access to the Oracle source database during the migration process. This capability utilizes Oracle read-only tables and partitions, which create a consistent view of the data during replication. Additionally, you can replicate data from an on-premises or EC2 instance database to an Aurora MySQL database using the Amazon DMS replication instance, creating an Aurora MySQL replica.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	MySQL doesn't support the READ ONLY, you can use a workaround.

Oracle usage

Beginning with Oracle 11g, tables can be marked as read-only to prevent DML operations from altering table data.

Prior to Oracle 11g, the only way to set a table to read-only mode was by limiting table privileges to SELECT. The table owner was still able to perform read and write operations. Starting from Oracle 11g, users can run an ALTER TABLE statement and change the table mode to either READ ONLY or READ WRITE.

Oracle 12c Release 2 introduces greater granularity for read-only objects and supports read-only table partitions. Any attempt to perform a DML operation on a partition, or sub-partition, set to READ ONLY results in an error.

SELECT FOR UPDATE statements aren't allowed.

DDL operations are permitted if they don't modify table data.

Operations on indexes are allowed on tables set to READ ONLY mode.

Examples

```
CREATE TABLE EMP_READ_ONLY (  
EMP_ID NUMBER PRIMARY KEY,  
EMP_FULL_NAME VARCHAR2(60) NOT NULL);  
  
INSERT INTO EMP_READ_ONLY VALUES(1, 'John Smith');  
  
1 row created  
  
ALTER TABLE EMP_READ_ONLY READ ONLY;  
  
INSERT INTO EMP_READ_ONLY VALUES(2, 'Steven King');  
  
ORA-12081: update operation not allowed on table "SCT"."TBL_READ_ONLY"  
  
ALTER TABLE EMP_READ_ONLY READ WRITE;  
  
INSERT INTO EMP_READ_ONLY VALUES(2, 'Steven King');  
  
1 row created  
  
COMMIT;  
  
SELECT * FROM EMP_READ_ONLY;  
  
EMP_ID  EMP_FULL_NAME  
1       John Smith  
2       Steven King
```

For more information, see [ALTER TABLE](#) and [Changes in This Release for Oracle Database VLDB and Partitioning Guide](#) in the *Oracle documentation*.

MySQL usage

MySQL doesn't provide a built-in feature for read only tables, but the same functionality can be achieved using Aurora Replicas. The main disadvantage of this approach is that you must use two separated instances.

It is important to note that there is a granularity difference between this workaround and options with Oracle. you cannot mimic a single read-only table, this workaround creates a read-only copy of the database.

Example

The following walkthrough demonstrates how to create an Aurora replica:

1. Sign in to the Amazon Management Console and choose **RDS**.
2. Choose **Instance actions** and choose **Create Aurora replica**.
3. Enter all required details and choose **Create**.
4. View the new record on the instances page. Make sure that the **Status** changes to **available** and the **Replication role** changes to **reader**.

For more information, see [Create an Amazon Aurora Read Replica from an RDS MySQL DB Instance](#) in the *Amazon Web Services News Blog*.

Oracle and MySQL table constraints

With Amazon DMS, you can enforce data integrity rules on tables in Oracle and MySQL databases during migration. Table constraints are database objects that define rules for the data in a table. They prevent invalid data from being entered into the database and maintain consistency across related tables. The following sections will provide details on supported constraint types, configuration options, and best practices for managing table constraints during database migration using Amazon DMS.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Constraints	MySQL doesn't support REF, ENABLE, DEFERRABLE, and DISABLE. Also, MySQL doesn't support constraints on views.

Oracle usage

Oracle provides six types of constraints to enforce data integrity on table columns. Constraints ensure data inserted into tables is controlled and satisfies logical requirements.

Oracle integrity constraint types

- **Primary key** — Enforces that row values in a specific column are unique and not null.
- **Foreign key** — Enforces that values in the current table exist in the referenced table.
- **Unique** — Prevents data duplication on a column, or combination of columns, and allows one null value.
- **Check** — Enforces that values comply with a specific condition.
- **Not null** — Enforces that null values can't be inserted into a specific column.
- **REF** — References an object in another object type or in a relational table.

Oracle constraint creation

You can create new constraints in two ways.

- **Inline** — Defines a constraint as part of a table column declaration.

```
CREATE TABLE EMPLOYEES (  
  EMP_ID NUMBER PRIMARY KEY,...);
```

- **Out-of-line** — Defines a constraint as part of the table DDL during table creation.

```
CREATE TABLE EMPLOYEES (EMP_ID NUMBER,...,  
  CONSTRAINT PK_EMP_ID PRIMARY KEY(EMP_ID));
```

Note

Declare NOT NULL constraints using the inline method.

Use the following syntax to specify Oracle constraints:

- CREATE / ALTER TABLE

- CREATE / ALTER VIEW

Note

Views have only a primary key, foreign key, and unique constraints.

Privileges

You need privileges on the table where constraints are created and, in case of foreign key constraints, you need the REFERENCES privilege on the referenced table.

PRIMARY KEY constraints

A unique identifier for each record in a database table can appear only once and can't contain NULL values. A table can only have one primary key.

When you create a primary key constraint inline, you can specify only the PRIMARY KEY keyword. When you create the constraint out-of-line, you must specify one column or a combination of columns.

Creating a new primary key constraint also implicitly creates a unique index on the primary key column if no index already exists. When dropping a primary key constraint, the system-generated index is also dropped. If a user defined index was used, the index isn't dropped.

- Primary keys can't be created on columns defined with the following data types: LOB, LONG, LONG RAW, VARRAY, NESTED TABLE, BFILE, REF, TIMESTAMP WITH TIME ZONE.

You can use the TIMESTAMP WITH LOCAL TIME ZONE data type as a primary key.

- Primary keys can be created from multiple columns (composite PK). They are limited to a total of 32 columns.
- Defining the same column as both a primary key and as a unique constraint isn't allowed.

Examples

Create an inline primary key using a system-generated primary key constraint name.

```
CREATE TABLE EMPLOYEES (
```

```
EMPLOYEE_ID NUMBER PRIMARY KEY,  
FIRST_NAME VARCHAR2(20),  
LAST_NAME VARCHAR2(25),  
EMAIL VARCHAR2(25));
```

Create an inline primary key using a user-specified primary key constraint name.

```
CREATE TABLE EMPLOYEES (  
  EMPLOYEE_ID NUMBER CONSTRAINT PK_EMP_ID PRIMARY KEY,  
  FIRST_NAME VARCHAR2(20),  
  LAST_NAME VARCHAR2(25),  
  EMAIL VARCHAR2(25));
```

Create an out-of-line primary key.

```
CREATE TABLE EMPLOYEES(  
  EMPLOYEE_ID NUMBER,  
  FIRST_NAME VARCHAR2(20),  
  LAST_NAME VARCHAR2(25),  
  EMAIL VARCHAR2(25));  
CONSTRAINT PK_EMP_ID PRIMARY KEY (EMPLOYEE_ID));
```

Add a primary key to an existing table.

```
ALTER TABLE SYSTEM_EVENTS  
  ADD CONSTRAINT PK_EMP_ID PRIMARY KEY (EVENT_CODE, EVENT_TIME);
```

FOREIGN KEY constraints

Foreign key constraints identify the relationship between column records defined with a foreign key constraint and a referenced primary key or a unique column. The main purpose of a foreign key is to enforce that the values in table A also exist in table B as referenced by the foreign key.

A referenced table is known as a parent table. The table on which the foreign key was created is known as a child table. Foreign keys created in child tables generally reference a primary key constraint in a parent table.

Limitations

Foreign keys can't be created on columns defined with the following data types: LOB, LONG, LONG RAW, VARRAY, NESTED TABLE, BFILE, REF, TIMESTAMP WITH TIME ZONE.

Composite foreign key constraints comprised from multiple columns can't have more than 32 columns.

Foreign key constraints can't be created in a CREATE TABLE statement with a subquery clause.

A referenced primary key or unique constraint on a parent table must be created before the foreign key creation command.

ON DELETE clause

The ON DELETE clause specifies the effect of deleting values from a parent table on the referenced records of a child table. If the ON DELETE clause isn't specified, Oracle doesn't allow deletion of referenced key values in a parent table that has dependent rows in the child table.

- ON DELETE CASCADE — Dependent foreign key values in a child table are removed along with the referenced values from the parent table.
- ON DELETE NULL — Dependent foreign key values in a child table are updated to NULL.

Examples

Create an inline foreign key with a user-defined constraint name.

```
CREATE TABLE EMPLOYEES (  
  EMPLOYEE_ID NUMBER PRIMARY KEY,  
  FIRST_NAME VARCHAR2(20),  
  LAST_NAME VARCHAR2(25),  
  EMAIL VARCHAR2(25) ,  
  DEPARTMENT_ID REFERENCES DEPARTMENTS(DEPARTMENT_ID));
```

Create an out-of-line foreign key with a system-generated constraint name.

```
CREATE TABLE EMPLOYEES (  
  EMPLOYEE_ID NUMBER PRIMARY KEY,  
  FIRST_NAME VARCHAR2(20),  
  LAST_NAME VARCHAR2(25),  
  EMAIL VARCHAR2(25),  
  DEPARTMENT_ID NUMBER,  
  CONSTRAINT FK_FEP_ID  
  FOREIGN KEY(DEPARTMENT_ID) REFERENCES DEPARTMENTS(DEPARTMENT_ID));
```

Create a foreign key using the ON DELETE CASCADE clause.

```
CREATE TABLE EMPLOYEES (  
  EMPLOYEE_ID NUMBER PRIMARY KEY,  
  FIRST_NAME VARCHAR2(20),  
  LAST_NAME VARCHAR2(25),  
  EMAIL VARCHAR2(25),  
  DEPARTMENT_ID NUMBER,  
  CONSTRAINT FK_FEP_ID  
  FOREIGN KEY(DEPARTMENT_ID) REFERENCES DEPARTMENTS(DEPARTMENT_ID)  
  ON DELETE CASCADE);
```

Add a foreign key to an existing table.

```
ALTER TABLE EMPLOYEES  
  ADD CONSTRAINT FK_FEP_ID  
  FOREIGN KEY(DEPARTMENT_ID) REFERENCES DEPARTMENTS(DEPARTMENT_ID);
```

UNIQUE constraints

A unique constraint is similar to a primary key constraint. It specifies that the values in a single column, or combination of columns, must be unique and can't repeat in multiple rows.

The main difference from primary key constraint is that a unique constraint can contain NULL values. NULL values in multiple rows are also supported provided the combination of values is unique.

Limitations

A unique constraint can't be created on columns defined with the following data types: LOB, LONG, LONG RAW, VARRAY, NESTED TABLE, BFILE, REF, TIMESTAMP WITH TIME ZONE.

A unique constraint comprised from multiple columns can't have more than 32 columns.

Primary key and unique constraints can't be created on the same column or columns.

Example

Create an inline unique Constraint.

```
CREATE TABLE EMPLOYEES (  
  EMPLOYEE_ID NUMBER PRIMARY KEY,
```

```
FIRST_NAME VARCHAR2(20),
LAST_NAME VARCHAR2(25),
EMAIL VARCHAR2(25) CONSTRAINT UNIQ_EMP_EMAIL UNIQUE,
DEPARTMENT_ID NUMBER);
```

Check constraints

Check constraints are used to validate values in specific columns that meet specific criteria or conditions. For example, you can use a check constraint on an `EMPLOYEE_EMAIL` column to validate that each record has an `@aws.com` suffix. If a record fails the check validation, an error is raised and the record isn't inserted.

Using a check constraint can help transfer some of the logical integrity validation from the application to the database.

When you create a check constraint as inline, it can only be defined on a specific column. When using the out-of-line method, the check constraint can be defined on multiple columns.

Limitations

Check constraints can't perform validation on columns of other tables.

Check constraints can't be used with functions that aren't deterministic (for example, `CURRENT_DATE`).

Check constraints can't be used with user-defined functions.

Check constraints can't be used with pseudo columns such as: `CURRVAL`, `NEXTVAL`, `LEVEL`, or `ROWNUM`.

Example

Create an inline check constraint that uses a regular expression to validate the email suffix of inserted rows contains `@aws.com`.

```
CREATE TABLE EMPLOYEES (
  EMPLOYEE_ID NUMBER PRIMARY KEY,
  FIRST_NAME VARCHAR2(20),
  LAST_NAME VARCHAR2(25),
  EMAIL VARCHAR2(25)
  CHECK(REGEXP_LIKE (EMAIL, '^[A-Za-z]+@aws.com?{1,3}$')),
  DEPARTMENT_ID NUMBER);
```

NOT NULL constraints

A NOT NULL constraint prevents a column from containing any null values. To enable the NOT NULL constraint, make sure that you specify the NOT NULL keyword during table creation (inline only). Permitting null values is the default if NOT NULL isn't specified.

Example

```
CREATE TABLE EMPLOYEES (  
  EMPLOYEE_ID NUMBER PRIMARY KEY,  
  FIRST_NAME VARCHAR2(20) NOT NULL,  
  LAST_NAME VARCHAR2(25) NOT NULL,  
  EMAIL VARCHAR2(25),  
  DEPARTMENT_ID NUMBER);
```

Referential constraints

Referential constraints define a relationship between a column of type REF and the object it references. The REF constraint can be created both inline and out-of-line. Both methods permit defining a scope constraint, a row identifier constraint, or a referential integrity constraint based on the REF column.

Examples

Create a new Oracle type object.

```
CREATE TYPE DEP_TYPE AS OBJECT (  
  DEP_NAME VARCHAR2(60),  
  DEP_ADDRESS VARCHAR2(300));
```

Create a table based on the previously created type object.

```
CREATE TABLE DEPARTMENTS_OBJ_T OF DEP_TYPE;
```

Create the EMPLOYEES table with a reference to the previously created DEPARTMENTS table that is based on the DEP_TYPE object:

```
CREATE TABLE EMPLOYEES (  
  EMP_NAME VARCHAR2(60),  
  EMP_EMAIL VARCHAR2(60),
```

```
EMP_DEPT REF DEPARTMENT_TYP REFERENCES DEPARTMENTS_OBJ_T);
```

Special constraint states

Oracle provides granular control of database constraint enforcement. For example, you can disable constraints temporarily while making modifications to table data.

Constraint states can be defined using the `CREATE TABLE` or `ALTER TABLE` statements. The following constraint states are supported:

- `DEFERRABLE` — Enables the use of the `SET CONSTRAINT` clause in subsequent transactions until a `COMMIT` statement is submitted.
- `NOT DEFERRABLE` — Disables the use of the `SET CONSTRAINT` clause.
- `INITIALLY IMMEDIATE` — Checks the constraint at the end of each subsequent SQL statement (this state is the default).
- `INITIALLY DEFERRED` — Checks the constraint at the end of subsequent transactions.
- `VALIDATE` or `NO VALIDATE` — These parameters depend on whether the constraint is `ENABLED` or `DISABLED`.
- `ENABLE` or `DISABLE` — Specifies if the constraint should be enforced after creation (`ENABLE` by default). Several options are available when using `ENABLE` or `DISABLE`:
 - `ENABLE VALIDATE` — Enforces that the constraint applies to all existing and new data.
 - `ENABLE NOVALIDATE` — Only new data complies with the constraint.
 - `DISABLE VALIDATE` — A valid constraint is created in disabled mode with no index.
 - `DISABLE NOVALIDATE` — The constraint is created in disabled mode without validation of new or existing data.

Examples

Create a unique constraint with a state of `DEFERRABLE`.

```
CREATE TABLE EMPLOYEES (  
  EMPLOYEE_ID NUMBER PRIMARY KEY,  
  FIRST_NAME VARCHAR2(20),  
  LAST_NAME VARCHAR2(25),  
  EMAIL VARCHAR2(25) CONSTRAINT UNIQ_EMP_EMAIL UNIQUE DEFERRABLE,  
  DEPARTMENT_ID NUMBER);
```

Modify the state of the constraint to ENABLE NOVALIDATE.

```
ALTER TABLE EMPLOYEES
  ADD CONSTRAINT CHK_EMP_NAME CHECK(FIRST_NAME LIKE 'a%')
  ENABLE NOVALIDATE;
```

Using existing indexes to enforce constraint integrity

Primary key and unique constraints can be created based on an existing index to enforce the constraint integrity instead of implicitly creating a new index during constraint creation.

Example

Create a unique constraint based on an existing index.

```
CREATE UNIQUE INDEX IDX_EMP_ID ON EMPLOYEES(EMPLOYEE_ID);

ALTER TABLE EMPLOYEES
  ADD CONSTRAINT PK_CON_UNIQ
  PRIMARY KEY(EMPLOYEE_ID) USING INDEX IDX_EMP_ID;
```

For more information, see [CREATE TABLE](#) in the *Oracle documentation*.

MySQL usage

MySQL supports the following types of table constraints:

- PRIMARY KEY
- FOREIGN KEY
- UNIQUE
- NOT NULL
- ENUM (unique to MySQL)
- SET (unique to MySQL)

Note

MySQL doesn't support Oracle REF constraint.

Similar to constraint declaration in Oracle, in MySQL you can create constraints in-line or out-of-line when you specify table columns.

You can specify MySQL constraints using CREATE or ALTER TABLE. Views aren't supported.

You need privileges on the table in which constraints are created. For foreign key constraints, you need the REFERENCES privilege.

Primary key constraints

Primary key constraints uniquely identify each record and can't contain a NULL value.

Primary key constraint marks the column on which the table's heap is sorted (in the InnoDB storage engine, like Oracle IOT).

Primary key constraint uses the same ANSI SQL syntax as Oracle.

You can create a primary key constraint on a single column, or on multiple columns (composite primary keys), as the only PRIMARY KEY in a table.

Primary key constraint creates a unique B-tree index automatically on the column, or group of columns, marked as the primary key of the table.

Constraint names can be generated automatically by MySQL. If a name is explicitly specified during constraint creation, the constraint name is PRIMARY.

Examples

Create an inline primary key constraint with a system-generated constraint name.

```
CREATE TABLE EMPLOYEES (  
  EMPLOYEE_ID NUMERIC PRIMARY KEY,  
  FIRST_NAME VARCHAR(20),  
  LAST_NAME VARCHAR(25),  
  EMAIL VARCHAR(25));
```

Create an out-of-line primary key constraint. For both examples, the constraint name is PRIMARY.

```
CREATE TABLE EMPLOYEES(  
  EMPLOYEE_ID NUMERIC,
```

```
FIRST_NAME VARCHAR(20),
LAST_NAME VARCHAR(25),
EMAIL VARCHAR(25),
CONSTRAINT PK_EMP_ID PRIMARY KEY (EMPLOYEE_ID));
```

or

```
CREATE TABLE EMPLOYEES(
  EMPLOYEE_ID NUMERIC,
  FIRST_NAME VARCHAR(20),
  LAST_NAME VARCHAR(25)
  EMAIL VARCHAR(25),
  CONSTRAINT PRIMARY KEY (EMPLOYEE_ID));
```

Add a primary key constraint to an existing table.

```
ALTER TABLE SYSTEM_EVENTS
  ADD CONSTRAINT PK_EMP_ID PRIMARY KEY (EVENT_CODE, EVENT_TIME);
```

or

```
ALTER TABLE SYSTEM_EVENTS
  ADD CONSTRAINT PRIMARY KEY (EVENT_CODE, EVENT_TIME);
```

or

```
ALTER TABLE SYSTEM_EVENTS
  ADD PRIMARY KEY (EVENT_CODE, EVENT_TIME);
```

Drop the primary key.

```
ALTER TABLE SYSTEM_EVENTS DROP PRIMARY KEY;
```

Foreign key constraints

Important notes about foreign key constraints:

- Enforces referential integrity in the database. Values in specific columns or group of columns must match the values from another table or column.
- Creating a FOREIGN KEY constraint in MySQL uses the same ANSI SQL syntax as Oracle.
- Can be created only out-of-line during table creation.

- Use the REFERENCES clause to specify the table referenced by the foreign key constraint.
- A table can have multiple FOREIGN KEY constraints to describe its relationships with other tables.
- Use the ON DELETE clause to handle cases of FOREIGN KEY parent records deletions such as cascading deletes.
- Use the ON UPDATE clause to handle cases of FOREIGN KEY parent records updates such as cascading updates.
- Foreign key constraint names are generated automatically by the database or specified explicitly during constraint creation.

ON DELETE clause

MySQL provides four options to handle cases where data is deleted from the parent table and a child table is referenced by a FOREIGN KEY constraint. By default, without specifying any additional options, MySQL uses the NO ACTION method and raises an error if the referencing rows still exist when the constraint is verified.

- ON DELETE CASCADE — Removes any dependent foreign key values in the child table along with the referenced values from the parent table.
- ON DELETE RESTRICT — Prevents the deletion of referenced values from the parent table and the deletion of dependent foreign key values in the child table.
- ON DELETE NO ACTION — Prevents the deletion of referenced values from the parent table and the deletion of dependent foreign key values in the child table (the same as RESTRICT).
- ON DELETE SET NULL — Deletes the row from the parent table and sets the foreign key column, or columns in the child table, to NULL. If you specify a SET NULL action, ensure you have not declared the columns in the child table as NOT NULL.

ON UPDATE clause

Handle updates on FOREIGN KEY columns is also available using the ON UPDATE clause, which shares the same options as the ON DELETE clause:

- ON UPDATE CASCADE
- ON UPDATE RESTRICT
- ON UPDATE NO ACTION

Note

Oracle doesn't provide an `ON UPDATE` clause.

Examples

Create an out-of-line foreign key constraint with a system-generated constraint name.

```
CREATE TABLE EMPLOYEES (  
    EMPLOYEE_ID NUMERIC PRIMARY KEY,  
    FIRST_NAME VARCHAR(20),  
    LAST_NAME VARCHAR(25),  
    EMAIL VARCHAR(25),  
    DEPARTMENT_ID NUMERIC,  
    CONSTRAINT FK_FEP_ID FOREIGN KEY(DEPARTMENT_ID)  
        REFERENCES DEPARTMENTS(DEPARTMENT_ID));
```

Create a foreign key using the `ON DELETE CASCADE` clause.

```
CREATE TABLE EMPLOYEES (  
    EMPLOYEE_ID NUMERIC PRIMARY KEY,  
    FIRST_NAME VARCHAR(20),  
    LAST_NAME VARCHAR(25),  
    EMAIL VARCHAR(25),  
    DEPARTMENT_ID NUMERIC,  
    CONSTRAINT FK_FEP_ID FOREIGN KEY(DEPARTMENT_ID)  
        REFERENCES DEPARTMENTS(DEPARTMENT_ID) ON DELETE CASCADE);
```

Add a foreign key to an existing table.

```
ALTER TABLE EMPLOYEES  
    ADD CONSTRAINT FK_FEP_ID  
        FOREIGN KEY(DEPARTMENT_ID)  
        REFERENCES DEPARTMENTS(DEPARTMENT_ID);
```

UNIQUE constraints

Important notes about unique constraints:

- Ensures that a value in a column, or a group of columns, is unique across the entire table.

- MySQL UNIQUE constraint syntax is ANSI SQL compatible.
- Automatically creates a B-tree index on the respective column, or a group of columns, when creating a UNIQUE constraint.
- If duplicate values exist in the column(s) on which the constraint was defined during UNIQUE constraint creation, the UNIQUE constraint creation fails and returns an error message.
- UNIQUE constraints in MySQL accept multiple NULL values, similar to Oracle.
- UNIQUE constraint naming can be system-generated or explicitly specified.

Example

Create an inline unique constraint ensuring uniqueness of values in the email column.

```
CREATE TABLE EMPLOYEES (  
    EMPLOYEE_ID NUMERIC PRIMARY KEY,  
    FIRST_NAME VARCHAR(20),  
    LAST_NAME VARCHAR(25),  
    EMAIL VARCHAR(25) UNIQUE,  
    DEPARTMENT_ID NUMERIC);
```

Disable integration check

In MySQL, you don't have an option to DISABLE the integration check, but there is a session variable for disabling checks at the session level.

The following example turns on integration checks in the session.

```
SET FOREIGN_KEY_CHECKS=1;
```

The following example turns off integration checks in the session.

```
SET FOREIGN_KEY_CHECKS=0;
```

Unique MySQL constraints

- ENUM — The value must be one of the values listed in the column definition or the internal numeric equivalent. The value can't be the error value. That is, 0 or the empty string. For a column defined as ENUM ('a' , 'b' , 'c'), the values such as ' ' , 'd' , or 'ax' are not valid and are rejected.

- SET — The value must be the empty string or a value consisting only of the values listed in the column definition separated by commas. For a column defined as SET(' a ' , ' b ' , ' c '), values such as ' d ' or ' a , b , c , d ' are not valid and are rejected.

Summary

Oracle constraint or parameter	MySQL constraint or parameter
PRIMARY KEY	PRIMARY KEY
NOT NULL	FOREIGN KEY
UNIQUE	UNIQUE
CHECK	Not supported, in some cases you can use ENUM and SET. Can be implemented with triggers.
NOT NULL	NOT NULL
REF	Not supported
DEFERRABLE	Not supported as keyword, you can use the FOREIGN_KEY_CHECKS parameter.
NOT DEFERRABLE	Not supported as keyword, you can use the FOREIGN_KEY_CHECKS parameter.
SET CONSTRAINTS	Not supported as keyword, you can use the FOREIGN_KEY_CHECKS parameter.
INITIALLY IMMEDIATE	Default, not supported as keyword.
INITIALLY DEFERRED	Not supported
ENABLE	Default, not supported as keyword.
DISBALE	Not supported as keyword, you can use the FOREIGN_KEY_CHECKS parameter.

Oracle constraint or parameter	MySQL constraint or parameter
ENABLE VALIDATE	Default, not supported as keyword
ENABLE NOVALIDATE	Not supported
DISABLE VALIDATE	Not supported
DISABLE NOVALIDATE	Default, not supported as keyword
USING_INDEX_CLAUSE	Not supported
View constraints	Not supported
Metadata: DBA_CONSTRAINTS	Metadata: TABLE_CONSTRAINTS .

For more information, see [How MySQL Deals with Constraints](#), and [FOREIGN KEY Constraints](#) in the *MySQL documentation*.

Oracle and MySQL temporary tables

Temporary tables are useful for storing intermediate results, handling large data sets, and improving query performance. The following sections will provide detailed instructions on migrating Oracle and MySQL temporary tables using Amazon DMS, ensuring a smooth transition to Amazon while preserving data and application behavior.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
			MySQL doesn't support GLOBAL temporary tables. MySQL can't read from multiple sessions. MySQL

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
			drops tables after the session ends.

Oracle usage

In Oracle, you can create temporary tables for storing data that exists only for the duration of a session or transaction.

Use the `CREATE GLOBAL TEMPORARY TABLE` statement to create a temporary table. This type of table has a persistent DDL structure, but not persistent data. It doesn't generate redo during DML.

Two of the primary use-cases for temporary tables include:

- Processing many rows as part of a batch operation while requiring staging tables to store intermediate results.
- Storing data required only for the duration of a specific session. When the session ends, the session data is cleared.

When using temporary tables, the data is visible only to the session that inserts the data into the table.

Oracle 18c introduces private temporary tables which are temporary tables that are only available during session or transaction. After session or transaction ends they are automatically dropped.

Oracle global temporary tables

Global temporary tables store data in the Oracle Temporary Tablespace.

DDL operations on a temporary table are permitted including `ALTER TABLE`, `DROP TABLE`, and `CREATE INDEX`.

Temporary tables can't be partitioned, clustered, or created as index-organized tables. Also, they don't support parallel `UPDATE`, `DELETE`, and `MERGE`.

Foreign key constraints can't be created on temporary tables.

Processing DML operations on a temporary table doesn't generate redo data. However, undo data for the rows and redo data for the undo data itself are generated.

Indexes can be created for a temporary table. They are treated as temporary indexes. Temporary tables also support triggers.

Temporary tables can't be named after an existing table object and can't be dropped while containing records, even from another session.

Session-specific and transaction-specific temporary table syntax

Use `ON COMMIT` to specifies whether the temporary table data persists for the duration of a transaction or a session.

Use `PRESERVE ROWS` when the session ends, all data is truncated but persists beyond the end of the transaction.

Use `DELETE ROWS` to truncate data after each commit. This is the default behavior.

Oracle 12c temporary table enhancements

Global temporary table statistics

Prior to Oracle 12c, statistics on temporary tables were common to all sessions. Oracle 12c introduces session-specific statistics for temporary tables. Statistics can be configured using the `DBMS_STATS` preference `GLOBAL_TEMP_TABLE_STATS`, which can be set to `SHARED` or `SESSION`.

Global temporary table undo

Performing DML operations on a temporary table doesn't generate Redo data, but does generate undo data that eventually, by itself, generates redo records. Oracle 12c provides an option to store the temporary undo data in the temporary tablespace itself. This feature is configured using the `temp_undo_enabled` parameter with the options `TRUE` or `FALSE`.

For more information, see [TEMP_UNDO_ENABLED](#) in the *Oracle documentation*.

Examples

Create an Oracle global temporary table with `ON COMMIT PRESERVE ROWS`.

```
CREATE GLOBAL TEMPORARY TABLE EMP_TEMP (
```

```
EMP_ID NUMBER PRIMARY KEY,  
EMP_FULL_NAME VARCHAR2(60) NOT NULL,  
AVG_SALARY NUMERIC NOT NULL)  
ON COMMIT PRESERVE ROWS;  
  
CREATE INDEX IDX_EMP_TEMP_FN ON EMP_TEMP(EMP_FULL_NAME);  
  
INSERT INTO EMP_TEMP VALUES(1, 'John Smith', '5000');  
  
COMMIT;  
  
SELECT * FROM SCT.EMP_TEMP;  
  
EMP_ID EMP_FULL_NAME AVG_SALARY  
1      John Smith      5000
```

Create an Oracle global temporary table with `ON COMMIT DELETE ROWS`.

```
CREATE GLOBAL TEMPORARY TABLE EMP_TEMP (  
  EMP_ID NUMBER PRIMARY KEY,  
  EMP_FULL_NAME VARCHAR2(60) NOT NULL,  
  AVG_SALARY NUMERIC NOT NULL)  
ON COMMIT DELETE ROWS;  
  
INSERT INTO EMP_TEMP VALUES(1, 'John Smith', '5000');  
  
COMMIT;  
  
SELECT * FROM SCT.EMP_TEMP;
```

For more information, see [CREATE TABLE](#) in the *Oracle documentation*.

MySQL usage

MySQL temporary tables share many similarities with Oracle global temporary tables. From a syntax perspective, MySQL temporary tables are referred to as temporary tables without global definition. The implementation is mostly identical.

In terms of differences, Oracle stores the temporary table structure (DDL) for repeated use — even after a database restart — but doesn't store rows persistently. MySQL implements temporary tables differently: the table structure (DDL) isn't stored in the database. When a session ends, the temporary table is dropped.

In MySQL, every session is required to create its own temporary tables. Each session can create its own private temporary tables, using identical table names.

In Oracle, the default behavior when the `ON COMMIT` clause is omitted is `ON COMMIT DELETE ROWS`. In MySQL, the default is `ON COMMIT PRESERVE ROWS` and you can't change it.

Note

In Amazon Relational Database Service (Amazon RDS) for MySQL version 8.0.13, the user-created temporary tables and internal temporary tables created by the optimizer are stored in session temporary tablespaces that are allocated to a session from a pool of temporary tablespaces. When a session disconnects its temporary tablespaces are truncated and released back to the pool. In previous releases temporary tables were created in the global temporary tablespace `ibtmp1` which did not return disk space to the operating system after temporary tables were dropped. The `innodb_temp_tablespaces_dir` variable defines the location where session temporary tablespaces are created. The default location is the `#innodb_temp` directory in the data directory. The `INNODB_SESSION_TEMP_TABLESPACES` table provides metadata about session temporary tablespaces. The global temporary tablespace `ibtmp1` now stores rollback segments for changes made to user-created temporary tables.

Example

```
CREATE TEMPORARY TABLE EMP_TEMP (
  EMP_ID INT PRIMARY KEY,
  EMP_FULL_NAME VARCHAR(60) NOT NULL,
  AVG_SALARY INT NOT NULL1;
```

Summary

Feature	Oracle	Aurora MySQL
Semantic	Global Temporary Table	Temporary Table
Create table	CREATE GLOBAL TEMPORARY...	CREATE TEMPORARY...

Feature	Oracle	Aurora MySQL
Accessible from multiple sessions	Yes	No
Temp table DDL persist after session end or database restart user-managed data files	Yes	No (dropped at the end of the session)
Create index support	Yes	Yes
Foreign key support	Yes	No
ON COMMIT default	COMMIT DELETE ROWS	ON COMMIT PRESERVE ROWS
ON COMMIT PRESERVE ROWS	Yes	Yes
ON COMMIT DELETE ROWS	Yes	No
Alter table support	Yes	Yes
Gather statistics	dbms_stats.gather_table_stats	ANALYZE
Oracle 12c GLOBAL_TEMP_TABLE_STATS	dbms_stats.set_table_prefs	ANALYZE

For more information, see [CREATE TEMPORARY TABLE Statement](#) in the *MySQL documentation*.

Oracle and MySQL triggers

Triggers are database objects that encapsulate procedural logic, facilitating data validation, auditing, and maintaining referential integrity constraints. System administrators, database developers, and data engineers may require triggers to enforce business rules, log data changes, or propagate updates across related tables. The following sections provide detailed guidance on creating, managing, and testing triggers within the context of Amazon DMS.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
			<p>MySQL doesn't support statement and system event triggers. Also, MySQL doesn't support CREATE OR REPLACE.</p>

Oracle usage

A trigger is a named program that is stored in the database and fired when a specified event occurs. The associated event causing a trigger to run can either be tied to a specific database table, database view, database schema, or the database itself.

Triggers can be run after:

- Data Manipulation Language (DML) statements such as DELETE, INSERT, or UPDATE.
- Data Definition Language (DDL) statements such as CREATE, ALTER, or DROP.
- Database events and operations such as SERVERERROR, LOGON, LOGOFF, STARTUP, or SHUTDOWN.

Trigger types

- **DML** triggers can be created on tables or views and fire when inserting, updating, or deleting data. Triggers can fire before or after DML command run.
- **INSTEAD OF** triggers can be created on a non-editable view. INSTEAD OF triggers provide an application-transparent method for modifying views that can't be modified by DML statements.
- **SYSTEM event** triggers are defined at the database or schema level including triggers that fire after specific events:
 - User log-on and log-off.
 - Database events such as startup or shutdown, DataGuard events, server errors.

Examples

Create a trigger that runs after a row is deleted from the PROJECTS table, or if the primary key of a project is updated.

```
CREATE OR REPLACE TRIGGER PROJECTS_SET_NULL
  AFTER DELETE OR UPDATE OF PROJECTNO ON PROJECTS
  FOR EACH ROW
  BEGIN
    IF UPDATING AND :OLD.PROJECTNO != :NEW.PROJECTNO OR DELETING THEN
      UPDATE EMP SET EMP.PROJECTNO = NULL
        WHERE EMP.PROJECTNO = :OLD.PROJECTNO;
    END IF;
  END;
/
```

Trigger created.

```
DELETE FROM PROJECTS WHERE PROJECTNO=123;

SELECT PROJECTNO FROM EMP WHERE PROJECTNO=123;

PROJECTNO
NULL
```

Create a SYSTEM or schema trigger on a table. The trigger fires if a DDL DROP command runs for an object in the HR schema. It prevents dropping the object and raises an application error.

```
CREATE OR REPLACE TRIGGER PREVENT_DROP_TRIGGER
  BEFORE DROP ON HR.SCHEMA
  BEGIN
    RAISE_APPLICATION_ERROR (num => -20000,
      msg => 'Cannot drop object');
  END;
/
```

Trigger created.

```
DROP TABLE HR.EMP

ERROR at line 1:
ORA-00604: error occurred at recursive SQL level 1
ORA-20000: Cannot drop object
```

```
ORA-06512: at line 2
```

For more information, see [CREATE TRIGGER Statement](#) in the *Oracle documentation*.

MySQL usage

MySQL supports triggers, but not all of the functionality provided by Oracle. Triggers are associated with users for privileges reasons and with specific tables. Triggers fire at the row level, and not at the statement level. You can modify MySQL triggers using a FOLLOWS or PRECEDES clause. Also, MySQL triggers can be chained using the FOLLOWS or PRECEDES clauses.

Syntax

```
CREATE
[DEFINER = { user | CURRENT_USER }]
TRIGGER trigger_name
trigger_time trigger_event
ON tbl_name FOR EACH ROW
[trigger_order]
trigger_body
trigger_time: { BEFORE | AFTER }
trigger_event: { INSERT | UPDATE | DELETE }
trigger_order: { FOLLOWS | PRECEDES } other_trigger_name
```

Examples

Create a trigger referencing the OLD and NEW values.

```
set delimiter /
CREATE OR REPLACE TRIGGER PROJECTS_SET_NULL
BEFORE UPDATE ON PROJECTS
FOR EACH ROW
BEGIN
IF OLD.PROJECTNO != NEW.PROJECTNO THEN
UPDATE EMP SET EMP.PROJECTNO = NULL
WHERE EMP.PROJECTNO = OLD.PROJECTNO;
END IF;
END;
/
set delimiter ;
UPDATE PROJECTS WHERE PROJECTNO=123;
```

```
SELECT PROJECTNO FROM EMP WHERE PROJECTNO=123;
PROJECTNO
-----
NULL
```

Drop a trigger.

```
DROP TRIGGER PROJECTS_SET_NULL
```

Summary

Trigger	Oracle	MySQL
Before update trigger, row level	<pre>CREATE OR REPLACE TRIGGER check_update BEFORE UPDATE ON projects FOR EACH ROW BEGIN /*Trigger body*/ END; /</pre>	<pre>CCREATE TRIGGER check_update BEFORE UPDATE ON projects FOR EACH ROW BEGIN /*Trigger body*/ END; /</pre>
Before update trigger, statement level	<pre>CREATE OR REPLACE TRIGGER check_update BEFORE UPDATE ON projects BEGIN /*Trigger body*/ END; /</pre>	Not supported
System or event trigger	<pre>CREATE OR REPLACE TRIGGER drop_trigger BEFORE DROP ON hr.SCHEMA BEGIN RAISE_APPLICATION_ER ROR (num => -20000,</pre>	Not supported

Trigger	Oracle	MySQL
	<pre> msg => 'Cannot drop object'); END; / </pre>	
Referencing :old and :new values in triggers	<p>Use :NEW and :OLD in trigger body:</p> <pre> CREATE OR REPLACE TRIGGER Upper-New DeleteOld BEFORE INSERT OR UPDATE OF first_name ON employees FOR EACH ROW BEGIN :NEW.first_name := UPPER(:NEW.first_n ame); :NE W.salary := :OLD.sala ry; END; / </pre>	<p>Use NEW and OLD in trigger body:</p> <pre> CREATE TRIGGER UpperNewD eleteOld BEFORE UPDATE ON empys FOR EACH ROW SET NEW.first_name = UPPER(NEW.first_na me), NEW.salary = OLD.salar y; END; / </pre>
Database event level trigger	<pre> CREATE TRIGGER register_ shutdown ON DATABASE SHUTDOWN BEGIN Insert into logging values ('DB was shut down', sysdate); commit; END; / </pre>	Not supported

Trigger	Oracle	MySQL
Drop a trigger	<pre data-bbox="597 226 1019 338">DROP TRIGGER last_name _change_trg;</pre>	<pre data-bbox="1073 226 1495 338">DROP TRIGGER last_name _change_trg;</pre>
Modify logic run by a trigger	<p data-bbox="597 380 1019 464">Can be used with create or replace</p> <pre data-bbox="597 506 1019 1010">CREATE OR REPLACE TRIGGER UpperNewDeleteOld BEFORE INSERT OR UPDATE OF first_name ON employees FOR EACH ROW BEGIN <<NEW CONTENT>> END; /</pre>	<p data-bbox="1073 380 1495 422">Not supported</p>
Enable a trigger	<pre data-bbox="597 1058 1019 1209">ALTER TRIGGER UpperNewD eleteOld ENABLE;</pre>	<p data-bbox="1073 1058 1495 1335">Not supported. Can be achieved by setting variables for each trigger to determine if it is turned off or turned on, and then checking the variable in an IF statement.</p>
Disable a trigger	<pre data-bbox="597 1377 1019 1528">ALTER TRIGGER UpperNewD eleteOld DISABLE;</pre>	<p data-bbox="1073 1377 1495 1654">Not supported. Can be achieved by setting variables for each trigger to determine if it is turned off or turned on, and then checking the variable in an IF statement.</p>

For more information, see [Trigger Syntax and Examples](#) and [CREATE TRIGGER Statement](#) in the *MySQL documentation*.

Oracle and MySQL tablespaces and data files

Oracle and MySQL databases use tablespaces and data files to store data. A tablespace is a logical storage unit, while a data file is a physical file that stores data for the tablespace. The following sections provide details on working with tablespaces and data files during database migration.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Aurora MySQL doesn't support tablespace for each file only and physical files attributes.

Oracle usage

The storage structure of an Oracle database contains both physical and logical elements.

- **Tablespaces** — Each Oracle database contains one or more tablespaces, which are logical storage groups used as containers for creating new tables and indexes.
- **Data files** — Each tablespace is made up of one or more data files, which are the physical elements of an Oracle database tablespace. Datafiles can be located on the local file system, located in raw partitions, managed by Oracle ASM, or located on a network file system.

Storage hierarchy

- **Database** — Each Oracle database is composed of one or more tablespaces.
- **Tablespace** — Each Oracle tablespace is composed of one or more data files. Tablespaces are logical entities that have no physical manifestation on the file system.
- **Data files** — Physical files located on a file system. Each Oracle tablespace consists of one or more data files.
- **Segments** — Each segment represents a single database object that consumes storage such as tables, indexes, and undo segments.

- **Extent** — Each segment consists of one or more extents. Oracle uses extents to allocate contiguous sets of database blocks on disk.
- **Block** — The smallest unit of I/O for reads and writes. For blocks storing table data, each block can store one or more table rows.

Types of Oracle database tablespaces

- **Permanent tablespaces** — Designated to store persistent schema objects for applications.
- **Undo tablespace** — A special type of system permanent tablespace used by Oracle to manage UNDO data when running the database in automatic undo management mode.
- **Temporary tablespace** — Contains schema objects valid for the duration of a session. It is also used for sort operations that can't fit into memory.

Tablespace privileges

Make sure that you meet the following criteria when you create a tablespace:

- The database user has the CREATE TABLESPACE system privilege.
- The database is in OPEN mode.

Examples

Create a USERS tablespace comprised of a single data file.

```
CREATE TABLESPACE USERS
  DATAFILE '/u01/app/oracle/oradata/orcl/users01.dbf' SIZE 5242880
  AUTOEXTEND ON NEXT 1310720 MAXSIZE 32767M
  LOGGING ONLINE PERMANENT BLOCKSIZE 8192
  EXTENT MANAGEMENT LOCAL AUTOALLOCATE DEFAULT
  NOCOMPRESS SEGMENT SPACE MANAGEMENT AUTO;
```

Drop a tablespace.

```
DROP TABLESPACE USERS;
OR
DROP TABLESPACE USERS INCLUDING CONTENTS AND DATAFILES;
```

For more information, see [CREATE TABLESPACE](#), [file_specification](#), and [DROP TABLESPACE](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL logical storage structure is similar to Oracle. It uses tablespaces for storing database objects, but the General Tablespace isn't supported. Only InnoDB file-per-table is provided.

Note

Starting from Amazon Relational Database Service (Amazon RDS) for MySQL version 8, you can rename a general tablespace using the `ALTER TABLESPACE ... RENAME TO` syntax.

- **Tablespace** — the directory where data files are stored.
- **Data files** — file-system files that are placed inside a tablespace (directory) and are used to store database objects such as tables or indexes. Created automatically by MySQL. Similar to how Oracle-Managed-Files (OMF) behave.

The InnoDB file-per-table feature applies to each InnoDB table. Its indexes are stored in a separate `.ibd` data file. Each `.ibd` data file represents an individual tablespace.

Tablespaces

After you create an Amazon Aurora MySQL cluster, three system tablespaces are automatically provisioned. You can't modify or drop them. These tablespaces hold database metadata or provide temporary storage for sorting and calculations:

- `innodb_system`
- `innodb_temporary`
- `innodb_file_per_table_n`

One of the main advantages when using Aurora MySQL is the reduced complexity of storage management. You don't need to create tablespaces because Aurora MySQL uses a unique, self-managed shared storage architecture. Database administrators don't need to manage most storage aspects of databases.

Example

View all tablespaces.

```
SELECT * FROM INFORMATION_SCHEMA.FILES;
```

Summary

Feature	Oracle	Amazon Aurora MySQL
Tablespace	Exists as a logical object and made from one or more user-specified or system-generated data files.	Exists as a logical object and consists of one data file.
Data file	<p>Can be explicitly created and resized by the user. Oracle-Managed-Files (OMF) support automatically created data files.</p> <p>Each data file can contain one or more tables and/or indexes.</p>	<p>The behavior is more like Oracle Managed Files (OMF).</p> <ul style="list-style-type: none"> Created automatically in the directory assigned to the tablespace. A single data file stores information for a specific table or index. Multiple data files can exist for a table or index.
Create a new tablespace with system-managed data files	<pre>CREATE TABLESPACE sales_tbs DATAFILE SIZE 400M;</pre>	Not supported
Create a new tablespace with user-managed data files	<pre>CREATE TABLESPACE sales_tbs DATAFILE '/oradata/ sales01.dbf' SIZE 1M AUTOEXTEND ON NEXT 1M;</pre>	Not supported

Feature	Oracle	Amazon Aurora MySQL
Alter the size of a datafile	<pre>ALTER DATABASE DATAFILE '/oradata/-sales01 .dbf' RESIZE 100M;</pre>	<pre>ALTER TABLE EMPLOYEES FORCE;</pre> <p>Reclaims free space in the data file, which can reduce and tablespace size.</p>
Add a datafile to an existing tablespace	<pre>ALTER TABLESPACE sales_tbs ADD DATAFILE '/oradata/ sales02.dbf' SIZE 10M;</pre>	Not supported
Per-database tablespace	<p>Supported as part of the Oracle 12c Multi-Tenant architecture. You can create different dedicated tablespaces for different pluggable databases and set as the default tablespace for a PDB:</p> <pre>ALTER SESSION SET CONTAINER = 'sales'; CREATE TABLESPACE sales_tbs DATAFILE '/oradata/ sales01.dbf' SIZE 1M AUTOEXTEND ON NEXT 1M; ALTER DATABASE sales TABLESPACE sales_tds;</pre>	Not supported

Feature	Oracle	Amazon Aurora MySQL
Metadata tables	Data Dictionary tables are stored in the SYSTEM tablespace.	Data Dictionary tables are stored in the innodb_system tablespace.
Tablespace data encryption	<p>Supported</p> <ul style="list-style-type: none"> Supported using transparent data encryption. Encryption and decryption are handled seamlessly. Users don't have to modify the application to access the data. 	<p>Supported</p> <ul style="list-style-type: none"> Encrypt using keys managed through Amazon KMS. Encryption and decryption are handled seamlessly. Users doesn't have to modify the application to access the data. Enable encryption while deploying a new cluster with the Amazon Management Console or API operations.

For more information, see [Encrypting Amazon RDS resources](#) in the *Amazon Relational Database Service User Guide*.

Oracle user-defined types

With Amazon DMS, you can migrate Oracle user-defined types (UDTs) to compatible Amazon database services. Oracle UDTs are custom data types that extend the built-in scalar data types, allowing you to store complex data structures, such as objects and collections.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		User-Defined Types	Aurora MySQL doesn't support user-defined types.

Oracle usage

Oracle refers to user-defined types (UDTs) as `OBJECT TYPES`. These types are managed using PL/SQL. User-defined types enable the creation of application-dedicated, complex data types that are based on, and extend, the built-in Oracle data types.

The `CREATE TYPE` statement supports creation of the following types:

- Objects types
- Varying array or `varray` types
- Nested table types
- Incomplete types
- Additional types such as an SQLJ object type, which is a Java class mapped to SQL user-defined type

Examples

Create an Oracle Object Type to store an employee phone number.

```
CREATE OR REPLACE TYPE EMP_PHONE_NUM AS OBJECT (
  PHONE_NUM VARCHAR2(11));

CREATE TABLE EMPLOYEES (
  EMP_ID NUMBER PRIMARY KEY,
  EMP_PHONE EMP_PHONE_NUM NOT NULL);

INSERT INTO EMPLOYEES VALUES(1, EMP_PHONE_NUM('111-222-333'));
SELECT a.EMP_ID, a.EMP_PHONE.PHONE_NUM FROM EMPLOYEES a;
```

```
EMP_ID  EMP_PHONE.P
1       111-222-333
```

Create an Oracle object type as a collection of attributes for the employees table.

```
CREATE OR REPLACE TYPE EMP_ADDRESS AS OBJECT (
  STATE VARCHAR2(2),
  CITY VARCHAR2(20),
  STREET VARCHAR2(20),
  ZIP_CODE NUMBER);

CREATE TABLE EMPLOYEES (
  EMP_ID NUMBER PRIMARY KEY,
  EMP_NAME VARCHAR2(10) NOT NULL,
  EMP_ADDRESS EMP_ADDRESS NOT NULL);

INSERT INTO EMPLOYEES VALUES(1, 'John Smith',
  EMP_ADDRESS('AL', 'Gulf Shores', '3033 Joyce Street', '36542'));

SELECT a.EMP_ID, a.EMP_NAME, a.EMP_ADDRESS.STATE,
  a.EMP_ADDRESS.CITY, a.EMP_ADDRESS.STREET, a.EMP_ADDRESS.ZIP_CODE
FROM EMPLOYEES a;
```

EMP_ID	EMP_NAME	STATE	CITY	STREET	ZIP_CODE
1	John Smith	AL	Gulf Shores	3033 Joyce Street	36542

For more information, see [CREATE TYPE](#) and [CREATE TYPE BODY](#) in the *Oracle documentation*.

MySQL usage

Currently, Amazon Aurora MySQL doesn't provide a directly comparable alternative for user-defined types.

Oracle unused columns

With Amazon DMS, you can identify and manage unused columns in Oracle databases during database migration and replication tasks.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	MySQL doesn't support unused columns.

Oracle usage

Oracle provides a method to mark columns as *unused*. Unused columns aren't physically dropped, but are treated as if they were dropped. Unused columns can't be restored. Select statements don't retrieve data from columns marked as unused and aren't displayed when running a DESCRIBE table command.

The main advantage of setting a column to UNUSED is to reduce possible high database load when dropping a column from a large table. To overcome this issue, a column can be marked as unused and then be physically dropped later.

To set a column as unused, use the SET UNUSED clause.

Examples

```
ALTER TABLE EMPLOYEES SET UNUSED (COMMISSION_PCT);
ALTER TABLE EMPLOYEES SET UNUSED (JOB_ID, COMMISSION_PCT);
```

Display unused columns.

```
SELECT * FROM USER_UNUSED_COL_TABS;
```

```
TABLE_NAME  COUNT
EMPLOYEES   3
```

Drop the column permanently (physically drop the column).

```
ALTER TABLE EMPLOYEES DROP UNUSED COLUMNS;
```

For more information, see [CREATE TABLE](#) in the *Oracle documentation*.

MySQL usage

Currently, Amazon Aurora MySQL doesn't provide a comparable alternative for unused columns.

Oracle virtual columns and MySQL generated columns

With Amazon DMS, you can seamlessly migrate databases that utilize virtual columns (Oracle) or generated columns (MySQL) to compatible target databases. Virtual columns and generated columns define values derived from other columns or expressions, providing a means to store computed data without modifying the base tables. This functionality is beneficial for applications relying on calculated fields, auditing requirements, or data denormalization strategies.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Creating Table	Different paradigm and syntax.

Oracle usage

Oracle virtual columns appear as normal columns, but their values are calculated instead of being stored in the database. You can't create virtual columns based on other virtual columns and can only reference columns from the same table. When you create a virtual column, you can either explicitly specify the data type or let the database select the data type based on the expression.

You can use virtual columns with constraints, indexes, table partitioning, and foreign keys.

Functions in expressions must be deterministic at the time of table creation.

Virtual columns can't be manipulated by DML operations.

You can use virtual columns in a WHERE clause and as part of DML commands.

When you create an index on a virtual column, Oracle creates a function-based index.

Virtual columns don't support index-organized tables, external, objects, clusters, or temporary tables.

The output of a virtual column expression must be a scalar value.

The virtual column keywords `GENERATED ALWAYS AS` and `VIRTUAL` aren't mandatory and are provided for clarity only.

```
COLUMN_NAME [data type] [GENERATED ALWAYS] AS (expression) [VIRTUAL]
```

The keyword `AS` after the column name indicates the column is created as a virtual column.

A virtual column doesn't need to be specified in an `INSERT` statement.

Examples

Create a table that includes two virtual columns.

```
CREATE TABLE EMPLOYEES (  
  EMPLOYEE_ID NUMBER,  
  FIRST_NAME VARCHAR2(20),  
  LAST_NAME VARCHAR2(25),  
  USER_NAME VARCHAR2(25),  
  EMAIL AS (LOWER(USER_NAME) || '@aws.com'),  
  HIRE_DATE DATE,  
  BASE_SALARY NUMBER,  
  SALES_COUNT NUMBER,  
  FINAL_SALARY NUMBER GENERATED ALWAYS AS  
    (CASE WHEN SALES_COUNT >= 10 THEN BASE_SALARY +  
      (BASE_SALARY * (SALES_COUNT * 0.05))  
    END)  
  VIRTUAL);
```

Insert a new record into the table without specifying values for the virtual column.

```
INSERT INTO EMPLOYEES  
  (EMPLOYEE_ID, FIRST_NAME, LAST_NAME,  
   USER_NAME, HIRE_DATE, BASE_SALARY, SALES_COUNT)  
VALUES(1, 'John', 'Smith', 'jsmith',  
      '17-JUN-2003', 5000, 21);
```

Select the email Virtual Column from the table.

```
SELECT email FROM EMPLOYEES;

EMAIL          FINAL_SALARY
jsmith@aws.com 10250
```

For more information, see [CREATE TABLE](#) in the *Oracle documentation*.

MySQL usage

The syntax and functionality of generated columns are similar to virtual columns. They appear as normal columns, but their values are calculated. Generated columns cannot be created based on other Generated Columns and can only reference columns from the same table. When you create generated columns, make sure that you explicitly specify the data type of the column.

- Unlike Oracle, you can create generated columns based on other generated columns preceding them in the field list.
- You can use generated columns with constraints, indexes, table partitioning.
- Functions in expressions must be deterministic at the time of table creation.
- Generated columns can't be manipulated by DML operations.
- Generated columns can be used in a WHERE clause and as part of DML commands.
- When you create an index on a generated column, the generated values are stored in the index.
- The output of a generated column expression must be a scalar value.

Examples

Create a table that includes two generated columns.

```
CREATE TABLE EMPLOYEES (
  EMPLOYEE_ID INT,
  FIRST_NAME VARCHAR(20),
  LAST_NAME VARCHAR(25),
  USER_NAME VARCHAR(25),
  EMAIL VARCHAR(25) AS
    (CONCAT(LOWER(USER_ NAME), '@aws.com')),
  HIRE_DATE DATE,
  BASE_SALARY INT,
  SALES_COUNT INT,
```

```
FINAL_SALARY INT GENERATED ALWAYS AS
(CASE WHEN SALES_COUNT >= 10 THEN BASE_SALARY + (
BASE_SALARY * (SALES_COUNT * 0.05)) END) VIRTUAL);
```

Insert a new record into the table without specifying values for the generated column.

```
INSERT INTO EMPLOYEES
(EMPLOYEE_ID, FIRST_NAME, LAST_NAME, USER_NAME, HIRE_DATE,
BASE_SALARY, SALES_COUNT)
VALUES(1, 'John', 'Smith', 'jsmith', now(), 5000, 21);
```

Select the email and the generated column from the table.

```
SELECT EMAIL, FINAL_SALARY FROM EMPLOYEES;
```

For the preceding example, the result looks as shown following.

email	FINAL_SALARY
jsmith@aws.com	10250

For more information, see [CREATE TABLE and Generated Columns](#) and [Secondary Indexes and Generated Columns](#) in the *MySQL documentation*.

MySQL overall indexes summary

MySQL supports multiple types of indexes using different indexing algorithms that can provide performance benefits for different types of queries.

Usage

The built-in MySQL index types include:

- **B-tree** — Default indexes that you can use for equality and range for the majority of queries. These indexes can operate against all data types. You can use B-tree indexes to retrieve NULL values. B-tree index values are sorted in ascending order by default.
- **Hash** — Hash Indexes are practical for equality operators. These types of indexes are rarely used because they aren't transaction-safe. This type of index is supported by MEMORY and NDB storage engines.

- **Full-text** — Full-text indexes are useful when the application needs to query large amount of text, using more complicated morphology attributes.
- **Spatial** — This index supports objects such as POINT and GEOMETRY to run geographic-related queries.

Note

Amazon Relational Database Service (Amazon RDS) for MySQL version supports descending indexes: DESC in an index definition is no longer ignored but causes storage of key values in descending order. Previously indexes could be scanned in reverse order but at a performance penalty. A descending index can be scanned in forward order which is more efficient. Descending indexes also make it possible for the optimizer to use multiple-column indexes when the most efficient scan order mixes ascending order for some columns and descending order for others. For more information, see [Descending Indexes](#) in the *MySQL documentation*.

CREATE INDEX synopsis

```
CREATE [UNIQUE | FULLTEXT | SPATIAL] INDEX index_name
  [index_type]
  ON tbl_name (key_part,...)
  [index_option]
  [algorithm_option | lock_option] ...
```

```
key_part:
  col_name [(length)] [ASC | DESC]
```

```
index_option:
  KEY_BLOCK_SIZE [=] value
  | index_type
  | WITH PARSEr parser_name
  | COMMENT 'string'
```

```
index_type:
  USING {BTREE | HASH}
```

```
algorithm_option:
  ALGORITHM [=] {DEFAULT | INPLACE | COPY}
```

```
lock_option:  
  LOCK [=] {DEFAULT | NONE | SHARED | EXCLUSIVE}
```

By default, the CREATE INDEX statement creates a B-tree index.

Examples

Oracle CREATE/DROP index.

```
CREATE UNIQUE INDEX IDX_EMP_ID ON EMPLOYEES (EMPLOYEE_ID DESC);  
DROP INDEX IDX_EMP_ID;
```

MySQL CREATE/DROP index.

```
CREATE UNIQUE INDEX IDX_EMP_ID ON EMPLOYEES (EMPLOYEE_ID DESC);  
DROP INDEX IDX_EMP_ID;
```

Oracle ALTER INDEX ... RENAME .

```
ALTER INDEX IDX_EMP_ID RENAME TO IDX_EMP_ID_OLD;
```

MySQL ALTER INDEX ... RENAME .

```
ALTER TABLE EMPLOYEES RENAME INDEX IDX_EMP_ID TO IDX_EMP_ID_OLD;
```

Oracle REBUILD INDEX.

```
ALTER INDEX IDX_EMP_ID REBUILD;
```

MySQL REINDEX (REBUILD) INDEX.

```
ANALYZE TABLE EMPLOYEES;
```

For more information, see [CREATE INDEX Statement](#), [ANALYZE TABLE Statement](#), and [ALTER TABLE Statement](#) in the *MySQL documentation*.

Summary

Oracle indexes types and features	MySQL compatibility	MySQL equivalent
B-tree Index	Supported	B-tree Index
Index-organized tables	Supported	Default behavior by InnoDB
Reverse key indexes	Not supported	N/A
Descending indexes	Supported	ABS (default) / DESC
B-tree cluster indexes	Not supported	N/A
Unique and non-unique indexes	Supported	Syntax is identical
Function-based indexes	Supported	Use generated columns
Application domain indexes	Not supported	N/A
BITMAP index or Bitmap join indexes	Not supported	N/A
Composite indexes	Supported	Multicolumn indexes
Invisible indexes	Not supported	N/A
Local and global indexes	Not supported	N/A
Partial indexes for partitioned tables (Oracle 12c)	Limited compatibility	Column prefix index
CREATE INDEX... or DROP INDEX...	Supported	High percentage of syntax similarity
ALTER INDEX... (general definitions)	Not supported	N/A

Oracle indexes types and features	MySQL compatibility	MySQL equivalent
ALTER INDEX... REBUILD	Supported	ANALYZE TABLE
ALTER INDEX... REBUILD ONLINE	Not supported	N/A
Index metadata	STATISTICS (Oracle USER_INDEXES)	<pre>SELECT DISTINCT TABLE_SCHEMA, TABLE_NAME, INDEX_NAME, INDEX_TYPE FROM INFORMATION_SCHEMA.STATISTICS;</pre>
Index tablespace allocation	Not supported	N/A
Index parallel operations	Not supported	N/A
Index compression	No direct equivalent to Oracle index key compression or advanced index compression	N/A

Oracle bitmap indexes

With Amazon DMS, you can optimize query performance for data warehousing and ad-hoc queries by creating Oracle bitmap indexes. You can use bitmap indexes to enhance the speed of complex queries involving conditions, joins, and aggregations on columns with a relatively small number of distinct values. Bitmap indexes can significantly improve query response times, especially for star schema queries common in data warehousing and business intelligence applications.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Indexes	MySQL doesn't support BITMAP index.

Oracle usage

Bitmap indexes are task-specific indexes best suited for providing fast data retrieval for OLAP workloads and are generally very fast for read-mostly scenarios. However, bitmap indexes don't perform well in heavy DML or OLTP workloads.

Unlike B-tree indexes where an index entry points to a specific table row, a bitmap index stores a bitmap for each index key.

Bitmap indexes are ideal for low-cardinality data filtering where the number of distinct values in a column is relatively small.

Example

Create an Oracle bitmap index.

```
CREATE BITMAP INDEX IDX_BITMAP_EMP_GEN ON EMPLOYEES(GENDER);
```

For more information, see [CREATE INDEX](#) in the *Oracle documentation*.

MySQL usage

Currently, Amazon Aurora MySQL doesn't provide a comparable alternative for bitmap indexes.

Oracle and MySQL B-tree indexes

With Amazon DMS, you can efficiently migrate data from your Oracle database to Aurora MySQL, while preserving existing B-tree indexes.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Indexes	N/A

Oracle usage

B-tree indexes (B stands for balanced), are the most common index type in a relational database and are used for a variety of common query performance enhancing tasks. You can define B-tree indexes as an ordered list of values divided into ranges. They provide superior performance by associating a key with a row or range of rows.

B-tree indexes contain two types of blocks: branch blocks for searching and leaf blocks for storing values. The branch blocks also contain the root branch, which points to lower-level index blocks in the B-tree index structure.

B-tree indexes are useful for primary keys and other high-cardinality columns. They provide excellent data access performance for a variety of query patterns such as exact match searches and range searches. B-tree indexes are the default when you create a new index.

Example

Create a B-tree index.

```
CREATE INDEX IDX_EVENT_ID ON SYSTEM_LOG(EVENT_ID);
```

For more information, see [CREATE INDEX](#) in the *Oracle documentation*.

MySQL usage

MySQL provides full support for B-tree indexes. Certain constraints created in MySQL such as primary keys or unique keys are stored in a B-tree index format. Similar to Oracle, B-tree indexes are the default for new indexes.

The query optimizer in MySQL can use B-tree indexes when handling equality and range queries on data. The MySQL optimizer considers using B-tree indexes to access data, especially when queries use one or more of the following operators: `>`, `>=`, `<`, `#`, `=`.

In addition, query elements such as `IN`, `BETWEEN`, `IS NULL`, or `IS NOT NULL` can also use B-tree indexes for faster data retrieval.

There are two types of indexes: * **Clustered index** — A reference as primary key. When a primary key is defined on a table, InnoDB uses it as the clustered index. It is highly recommended to specify a primary key for all tables. If there is no primary key, MySQL locates the first `UNIQUE` index where all columns are `NOT NULL` and are used as a clustered index. If there is no primary key or `UNIQUE` index to use, InnoDB internally generates a hidden clustered index named `GEN_CLUST_INDEX`. * **Secondary index**: All indexes that are not clustered indexes. Each index entry has a reference to the clustered index. If the clustered index is applied on long values, the secondary indexes consume more storage space.

Example

Create a B-tree Index.

```
CREATE INDEX IDX_EVENT_ID ON SYSTEM_LOG (EVENT_ID);  
  
or  
  
CREATE INDEX IDX_EVENT_ID ON SYSTEM_LOG (EVENT_ID) USING BTREE;
```

For more information, see [CREATE INDEX Statement](#) in the *MySQL documentation*.

Oracle composite indexes and MySQL multiple-column indexes

With Amazon DMS, you can optimize database performance by creating composite indexes in Oracle databases and multiple-column indexes in MySQL databases. A composite index (Oracle) or multiple-column index (MySQL) is a database index built from multiple columns in a table, allowing queries to be satisfied by utilizing the index entries alone.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Indexes	N/A

Oracle usage

An index created on multiple table columns is known as a multi-column, concatenated, or composite index. The main purpose of composite indexes is to improve the performance of data retrieval for SELECT statements when filtering on all, or some, of the composite index columns. When using composite indexes, it is beneficial to place the most restrictive columns at the first position of the index to improve query performance. Column placement order is crucial when using composite indexes because the most prevalent columns are accessed first.

Examples

Create a composite index on the HR.EMPLOYEES table.

```
CREATE INDEX IDX_EMP_COMPI ON
EMPLOYEES (FIRST_NAME, EMAIL, PHONE_NUMBER);
```

Drop a composite index.

```
DROP INDEX IDX_EMP_COMPI;
```

For more information, see [Composite Indexes](#) in the *Oracle documentation*.

MySQL usage

MySQL multiple-column indexes are similar to composite indexes in Oracle.

These indexes are beneficial when queries filter on all indexed columns, the first indexed column, the first two indexed columns, the first three indexed columns, and so on. When indexed columns are specified in the optimal order during index creation, a single multiple-column index can improve performance in scenarios where several queries access the same database table.

You can specify up to 16 columns when creating a multiple-column index.

Examples

Create a multiple-column index on the EMPLOYEES table.

```
CREATE INDEX IDX_EMP_COMPI
ON EMPLOYEES (FIRST_NAME, EMAIL, PHONE_NUMBER);
```

Drop a multiple-column index.

```
DROP INDEX IDX_EMP_COMPI;
```

For more information, see [Multiple-Column Indexes](#) in the *MySQL documentation*.

Oracle function-based indexes and MySQL indexing on generated columns

With Amazon DMS, you can improve query performance by creating indexes on computed values or expressions in your databases. Oracle function-based indexes and MySQL indexes on generated columns let you index data derived from an expression or function, rather than just indexing on a column's stored values. This can significantly speed up queries that filter or sort on calculated values.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Indexes	MySQL doesn't support functional indexes, a workaround is available.

Oracle usage

Function-based indexes allow functions to be used in the WHERE clause of queries on indexed columns. Function-based indexes store the output of a function applied on the values of a table

column. The Oracle query optimizer only uses a function-based index when the function is used as part of a query.

Oracle updates the index for each DML to ensure that the value that returns from the function is correct.

Example

Create a function-based index.

```
CREATE TABLE SYSTEM_EVENTS(  
  EVENT_ID NUMERIC PRIMARY KEY,  
  EVENT_CODE VARCHAR2(10) NOT NULL,  
  EVENT_DESCRIPTION VARCHAR2(200),  
  EVENT_TIME TIMESTAMPTO_TIMESTAMP;  
  
CREATE INDEX EVNT_BY_DAY ON SYSTEM_EVENTS(  
  EXTRACT(DAY FROM EVENT_TIME));
```

For more information, see [Indexes and Index-Organized Tables](#) and [CREATE INDEX](#) in the *Oracle documentation*.

MySQL usage

MySQL does not directly support a feature equivalent to Oracle function-based indexes. However, workarounds exist that can offer similar functionality. Specifically, you can create secondary indexes on MySQL generated columns. Implementing this workaround may require modification of existing SQL queries.

A generated column derives its values from the result of an expression. Creating an index on a generated column allows the generated column to be used in a `WHERE` clause of a query while accessing data with the index. Unlike Oracle function-based indexes, this workaround requires specifying the function in the table column specification.

You can create generated columns as `STORED` or `VIRTUAL`. For our purposes, we need to create generated columns as `STORED`. Otherwise, we won't be able to index those columns.

MySQL can't use stored routines or functions with generated columns.

Generated columns support `NOT NULL` restrictions.

A generated expression cannot exceed 64 KB for the entire table. For example, you can create a single field with a generated expression length of 64 KB or 12 fields with a length of 5 KB each.

A generated column can't refer to itself or to other generated columns defined later, but it can refer to any previously defined generated columns.

The generation expression can only call native deterministic functions.

You can mix VIRTUAL and STORED columns within a table.

When you insert data to the table, make sure that you don't reference the generated columns in your insert statement.

Examples

Create a generated column that calculates the yearly salary based on the monthly salary, and create a secondary index on that column.

```
CREATE TABLE EMPS (ID INT, MONTH_SALARY INT,  
    YEAR_SALARY INT GENERATED ALWAYS AS (MONTH_SALARY*12),  
    INDEX FBI_YEAR_IDX (YEAR_SALARY));
```

```
INSERT INTO EMPS (ID, MONTH_SALARY) VALUES (1,10000);  
INSERT INTO EMPS (ID, MONTH_SALARY) VALUES (2,8764);  
INSERT INTO EMPS (ID, MONTH_SALARY) VALUES (3,4355);  
INSERT INTO EMPS (ID, MONTH_SALARY) VALUES (4,6554);
```

```
SELECT * FROM EMPS;
```

ID	MONTH_SALARY	YEAR_SALARY
1	10000	120000
2	8764	105168

Queries can reference the YEAR_SALARY column as part of the WHERE clause and access data using the FBI_YEAR_IDX index.

```
SELECT * FROM EMPS WHERE YEAR_SALARY>80000;  
SELECT * FROM EMPS WHERE MONTH_SALARY*12>80000;
```

Consider another example.

Create two generated columns using string manipulation functions as part of the table specification with secondary indexes on each.

```
CREATE TABLE EMPS (ID INT, FULL_NAME CHAR(40),
FIRST_NAME CHAR(20) GENERATED ALWAYS AS
  (SUBSTRING(FULL_NAME, 1, INSTR(FULL_NAME, ' '))),
LAST_NAME CHAR(20) GENERATED ALWAYS AS
  (SUBSTRING(FULL_NAME, INSTR(FULL_NAME, ' '))),
INDEX FBI_FNAME_IDX (FIRST_NAME),
INDEX FBI_LNAME_IDX (LAST_NAME));

INSERT INTO EMPS (ID, FULL_NAME) VALUES (1, 'James Kirk');
INSERT INTO EMPS (ID, FULL_NAME) VALUES (2, 'Benjamin Sisko');
INSERT INTO EMPS (ID, FULL_NAME) VALUES (3, 'Karthryn Janeway');
INSERT INTO EMPS (ID, FULL_NAME) VALUES (4, 'Jean- Luc Picard');
```

Queries can now use the FBI_FNAME_IDX index.

```
SELECT ID FROM EMPS WHERE
  SUBSTRING(FULL_NAME, 1, INSTR(FULL_NAME, ' '))='Jacob';

SELECT ID FROM EMPS WHERE FIRST_NAME='Jacob';
```

Note

For the preceding example, generated columns were not necessary. However, the generated columns were provided as an example. Instead, you can use a B-tree index created on the column prefix to achieve the same results.

```
CREATE TABLE EMPS (ID INT, FULL_NAME CHAR(40));
CREATE INDEX FBI_NAME_PREF_IDX ON EMPS (FULL_NAME(20));
```

For more information, see [CREATE TABLE and Generated Columns](#) in the *MySQL documentation*.

Oracle and MySQL invisible indexes

With Amazon DMS, you can create and manage invisible indexes in Oracle and MySQL databases, providing a way to evaluate the potential benefits of an index before making it visible and impacting workload performance.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Indexes	MySQL doesn't support invisible indexes.

Oracle usage

In Oracle, the invisible index feature gives database administrators the ability to create indexes, or change existing indexes, that are ignored by the optimizer. They are maintained during DML operations and are kept relevant, but are different from usable indexes.

The most common use cases for invisible indexes are:

- Testing the effect of a dropped index without actually dropping it.
- Using a specific index for certain operations or modules of an application without affecting the overall application.
- Adding an index to a set of columns on which an index already exists.

Database administrators can force the optimizer to use invisible indexes by changing the `OPTIMIZER_USE_INVISIBLE_INDEXES` parameter to true. You can use invisible indexes if they are specified as a HINT.

Examples

Change an index to an invisible index.

```
ALTER INDEX idx_name INVISIBLE;
```

Change an invisible index to a visible index.

```
ALTER INDEX idx_name VISIBLE;
```

Create an invisible index.

```
CREATE INDEX idx_name ON employees(first_name) INVISIBLE;
```

Query all invisible indexes.

```
SELECT TABLE_OWNER, INDEX_NAME FROM DBA_INDEXES
WHERE VISIBILITY = 'INVISIBLE';
```

For more information, see [Understand When to Use Unusable or Invisible Indexes](#) in the *Oracle documentation*.

MySQL usage

Amazon Relational Database Service (Amazon RDS) for MySQL version 8 supports invisible indexes. An invisible index is not used by the optimizer at all but is otherwise maintained normally. Indexes are visible by default.

Invisible indexes make it possible to test the effect of removing an index on query performance without making a destructive change that must be undone should the index turn out to be required.

For more information, see [Invisible Indexes](#) in the *MySQL documentation*.

Oracle index-organized table and MySQL InnoDB clustered index

With Amazon DMS, you can migrate databases that utilize Oracle index-organized tables and MySQL InnoDB clustered indexes.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Indexes	MySQL doesn't support the index-organized tables. This is the default behavior for InnoDB.

Oracle usage

In Oracle, an index-organized table (IOT) object is a special type of index/table hybrid that physically controls how data is stored at the table and index level. When you create a common database table or a heap-organized table, the data is stored unsorted, as a heap. However, when you create an index-organized table, the actual table data is stored in a B-tree index structure sorted by the primary key of each row. Each leaf block in the index structure stores both the primary key and non-key columns.

IOTs provide performance improvements when accessing data using the primary key because table records are sorted or clustered using the primary key and physically co-located alongside the primary key.

Example

Create an Oracle index-organized table storing ordered data based on the primary key.

```
CREATE TABLE SYSTEM_EVENTS (  
  EVENT_ID NUMBER,  
  EVENT_CODE VARCHAR2(10) NOT NULL,  
  EVENT_DESCRIPTION VARCHAR2(200),  
  EVENT_TIME DATE NOT NULL,  
  CONSTRAINT PK_EVENT_ID PRIMARY KEY(EVENT_ID))  
  ORGANIZATION INDEX;  
  
INSERT INTO SYSTEM_EVENTS VALUES(9, 'EVNT-A1-10', 'Critical', '01-JAN-2017');  
INSERT INTO SYSTEM_EVENTS VALUES(1, 'EVNT-C1-09', 'Warning', '01-JAN-2017');  
INSERT INTO SYSTEM_EVENTS VALUES(7, 'EVNT-E1-14', 'Critical', '01-JAN-2017');  
  
SELECT * FROM SYSTEM_EVENTS;  
  
EVENT_ID  EVENT_CODE  EVENT_DESCRIPTION  EVENT_TIM  
1         EVNT-C1-09  Warning           01-JAN-17  
7         EVNT-E1-14  Critical          01-JAN-17  
9         EVNT-A1-10  Critical          01-JAN-17
```

Note

The records are sorted in the reverse order from which they were inserted.

For more information, see [Indexes and Index-Organized Tables](#) in the *Oracle documentation*.

MySQL usage

MySQL doesn't support index-organized tables. However it provides similar functionality using InnoDB, which is the Amazon Aurora default storage engine.

Each InnoDB table provides a special clustered index. When you create a PRIMARY KEY on a table, InnoDB automatically uses it as the clustered index. This behavior is similar to index-organized tables in Oracle.

The best practice is to specify a primary key for each MySQL table. If you do not specify a primary key, MySQL locates the first unique index where all key columns are specified as NOT NULL and uses it as the clustered index.

If a table layout doesn't logically provide a column or multiple columns that are unique and not null, it is recommended to explicitly add an auto-incremented column to generate unique values.

Note

If no primary key or a suitable unique index can be found, InnoDB actually creates a hidden GEN_CLUST_INDEX clustered index with internally generated row ID values. These auto-generated row IDs are based on a 6-byte field that increases monotonically.

Example

Create a new table with a simple primary key. Because the storage engine is InnoDB, the table is created as a clustered table sorting data based on the primary key itself.

```
CREATE TABLE SYSTEM_EVENTS (  
  EVENT_ID INT PRIMARY KEY,  
  EVENT_CODE VARCHAR(10) NOT NULL,  
  EVENT_DESCRIPTION VARCHAR(200),  
  EVENT_TIME DATE NOT NULL);  
  
INSERT INTO SYSTEM_EVENTS VALUES(9, 'EVNT10', 'Critical', NOW());  
INSERT INTO SYSTEM_EVENTS VALUES(1, 'EVNT09', 'Warning', NOW());  
INSERT INTO SYSTEM_EVENTS VALUES(7, 'EVNT14', 'Critical', NOW());  
  
SELECT * FROM SYSTEM_EVENTS;
```

event_id	event_code	event_description	event_time
1	EVNT-C1-09	Warning	2017-01-01
7	EVNT-E1-14	Critical	2017-01-01
9	EVNT-A1-10	Critical	2017-01-01

For more information, see [Clustered and Secondary Indexes](#) in the *MySQL documentation*.

Oracle local and global partitioned indexes and MySQL partitioned indexes

With Amazon DMS, you can migrate data from Oracle databases using local and global partitioned indexes and from MySQL databases using partitioned indexes. Partitioned indexes facilitate operations on large datasets by logically dividing data into smaller subsets, improving performance and manageability.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Indexes	N/A

Oracle usage

Local and global indexes are used for partitioned tables in Oracle databases. Each index created on a partitioned table can be specified as either local or global.

- **Local partitioned index** maintains a one-to-one relationship between the index partitions and the table partitions. For each table partition, Oracle creates a separate index partition. This type of index is created using the LOCAL clause. Because each index partition is independent, index maintenance operations are easier and can be performed independently. Local partitioned indexes are managed automatically by Oracle during creation or deletion of table partitions.
- **Global partitioned index** contains keys from multiple table partitions in a single index partition. This type of index is created using the GLOBAL clause during index creation. A global index can be partitioned or non-partitioned. The default option is non-partitioned. When you create

global partitioned indexes on partitioned tables, certain restrictions exist for index management and maintenance. For example, dropping a table partition causes the global index to become unusable without an index rebuild.

Examples

Create a local index on a partitioned table.

```
CREATE INDEX IDX_SYS_LOGS_LOC ON SYSTEM_LOGS (EVENT_DATE)
LOCAL
(PARTITION EVENT_DATE_1,
PARTITION EVENT_DATE_2,
PARTITION EVENT_DATE_3);
```

Create a global index on a partitioned table.

```
CREATE INDEX IDX_SYS_LOGS_GLOB ON SYSTEM_LOGS (EVENT_DATE)
GLOBAL PARTITION BY RANGE (EVENT_DATE) (
PARTITION EVENT_DATE_1 VALUES LESS THAN (TO_DATE('01/01/2015', 'DD/MM/YYYY')),
PARTITION EVENT_DATE_2 VALUES LESS THAN (TO_DATE('01/01/2016', 'DD/MM/YYYY')),
PARTITION EVENT_DATE_3 VALUES LESS THAN (TO_DATE('01/01/2017', 'DD/MM/YYYY')),
PARTITION EVENT_DATE_4 VALUES LESS THAN (MAXVALUE));
```

For more information, see [Partitioning Concepts](#) and [Index Partitioning](#) in the *Oracle documentation*.

MySQL usage

Indexes created on partitioned tables are similar to local indexes in Oracle. MySQL doesn't provide an equivalent for Oracle global indexes because in MySQL, partitioning applies to all data and indexes of a table. It is not possible to partition only the data and not the indexes. All indexes on partitioned tables behave like an Oracle local index.

Examples

Drop a partition (the index that is used without a rebuild). Note that the run plan shows the scanned partitions.

```
ALTER TABLE SYSTEM_LOGS add INDEX EVENT_NO_IDX (EVENT_NO);
```

```
EXPLAIN SELECT * from SYSTEM_LOGS where EVENT_NO=2;
```

```
id select_type table partitions type possible_keys key
key_len ref rows filtered
1 SIMPLE SYSTEM_LOGS warning,critical ref EVENT_NO_IDX EVENT_NO_IDX 4
const 1 100
```

```
ALTER TABLE SYSTEM_LOGS DROP PARTITION critical;
EXPLAIN SELECT * from SYSTEM_LOGS where EVENT_NO=2;
```

```
id select_type table partitions type possible_keys key key_len
ref rows filtered
1 SIMPLE SYSTEM_LOGS warning ref EVENT_NO_IDX EVENT_NO_IDX 4
const 1 100
```

For more information, see [Overview of Partitioning in MySQL](#) in the *MySQL documentation*.

Oracle automatic indexing

With Amazon DMS, you can leverage Oracle automatic indexing to optimize database performance and reduce manual tuning efforts.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Indexes	MySQL doesn't provide an automatic indexing feature.

Oracle usage

Oracle 19 introduces the automatic indexing feature. This feature automates the index management tasks by automatically creating, rebuilding, and dropping indexes based on the changes in application workload, thus improving database performance.

Important functionality provided by automatic indexing:

- Automatic indexing process runs in the background at a predefined time interval and analyzes application workload. It identifies the tables/columns that are candidates for new indexes and creates new indexes.
- The auto indexes are initially created as invisible indexes. These invisible auto indexes are verified against SQL statements and if the performance is improved, then these indexes are converted as visible indexes.
- Identify and drop any existing under-performing auto indexes or any auto indexes not used for long period.
- Rebuilds the auto indexes that are marked unusable due to DDL operations.
- Provides package `DBMS_AUTO_INDEX` to configure automatic indexing and for generating reports related to automatic indexing operations.

Note

Up-to-date table statistics are very important for the auto indexing to function efficiently. Tables without statistics or with stale statistics aren't considered for auto indexing.

Oracle uses the `DBMS_AUTO_INDEX` package to configure auto indexes and generating reports. Following are some of the configuration options which can be set by using `CONFIGURE` procedure of `DBMS_AUTO_INDEX` package:

- Turning on and turning off automatic indexing in a database.
- Specifying schemas and tables that can use auto indexes.
- Specifying a retention period for unused auto indexes. By default, the unused auto indexes are deleted after 373 days.
- Specifying a retention period for unused non-auto indexes.
- Specifying a tablespace and a percentage of tablespace to store auto indexes.

Following are some of the reports related to automatic indexing operations which you can generate using `REPORT_ACTIVITY` and `REPORT_LAST_ACTIVITY` functions of the `DBMS_AUTO_INDEX` package.

- Report of automatic indexing operations for a specific period.
- Report of the last automatic indexing operation.

For more information, see [Managing Indexes](#) in the *Oracle documentation*.

MySQL usage

Currently, Amazon Aurora MySQL doesn't provide a comparable alternative for automatic indexing. The most reasonable option would be to run a scheduled set of queries to estimate if additional indexes are needed.

The following queries can help determine that.

Find user-tables without primary keys.

```
SELECT tab.table_schema,tab.table_name
FROM information_schema.tables tab
LEFT JOIN information_schema.table_constraints tco
  ON tab.table_schema = tco.table_schema
  AND tab.table_name = tco.table_name
  AND tco.constraint_type = 'PRIMARY KEY'
WHERE tco.constraint_type is null
  AND tab.table_schema not in('information_schema', 'performance_schema', 'sys')
  AND tab.table_type = 'BASE TABLE'
ORDER BY tab.table_schema, tab.table_name;
```

Unused indexes that can probably be dropped.

```
SELECT * FROM sys.schema_unused_indexes;
```

All of these should not be implemented in a script to decide if indexes should be created or dropped in a production environment. The Oracle Automatic Indexes will first assess if a new index is needed and if so, it will create an invisible index and only after ensuring nothing is harmed, then the index will become visible. You can't use this process in MySQL to avoid any production performance issues.

Special Oracle features and future MySQL content

This section provides reference pages for special Oracle features and MySQL current equivalents, alternatives, and future features.

Topics

- [Oracle Advanced Queuing and MySQL integration with Lambda](#)
- [Oracle and MySQL character sets](#)
- [Oracle database links and MySQL fully-qualified table names](#)
- [Oracle DBMS_SCHEDULER and MySQL events](#)
- [Oracle external tables and MySQL integration with Amazon S3](#)
- [Oracle and MySQL inline views](#)
- [Oracle JSON document support and MySQL JSON](#)
- [Oracle materialized views and MySQL summary tables or views](#)
- [Oracle multitenant and MySQL databases](#)
- [Oracle Resource Manager and dedicated Amazon Aurora MySQL clusters](#)
- [Oracle SecureFile LOBs and MySQL large objects](#)
- [Oracle synonyms](#)
- [Oracle and MySQL views](#)
- [Oracle XML DB and MySQL XML](#)
- [Oracle table compression](#)
- [Oracle Log Miner and MySQL logs](#)
- [Oracle SQL Result Cache and MySQL Query Cache](#)

Oracle Advanced Queuing and MySQL integration with Lambda

With Amazon DMS, you can seamlessly migrate data from Oracle Advanced Queuing to Aurora MySQL, and integrate with Amazon Lambda for event-driven processing. Oracle Advanced Queuing provides message queuing capabilities for Oracle databases, while Amazon Lambda allows running code without provisioning or managing servers.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	Use Amazon Lambda and Amazon Simple Queue Service with Aurora MySQL.

Oracle usage

The Oracle Advanced Queuing (AQ) feature enables database-integrated message queuing functionality. It is based on Oracle Streams and optimizes data functions by storing messages, allocating the messages to different service queues, and transmitting the messages using Oracle Net Services, HTTP, and HTTPS. AQ is implemented using database tables.

Oracle provides the `oracle.jdbc.aq` Java package as an interface to AQ. It contains the following items:

- Classes:
 - `AQDequeueOptions` — Specifies the options for the dequeue operation.
 - `AQEnqueueOptions` — Specifies the options for the enqueue operation.
 - `AQFactory` — A factory class for AQ, which creates components such as agent or message properties.
 - `AQNotificationEvent` — New message notifications.
- Interfaces:
 - `AQAgent` — An identity of a user, producer, or consumer of a message.
 - `AQMessage` — An enqueued or dequeued message.
 - `AQMessageProperties` — Message properties such as:
 - Correlation
 - Sender
 - Delay
 - Expiration
 - Recipients

- Priority
- Ordering
- `AQNotificationListener` — A listener interface for receiving AQ notification events.
- `AQNotificationRegistration` — A registration to be notified when a new message is enqueued on a particular queue.

For more information, see [Introduction to Oracle Database Advanced Queuing](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL provides built-in integration with Lambda functions, which can be called from within the database and interact with Amazon Simple Notification Service (Amazon SNS). The integration with Lambda functions provides a powerful framework for using Amazon services to implement custom solutions with less code.

Examples

For examples, see [Amazon Simple Notification Service](#).

For more information, see [Invoking a Lambda function with an Aurora MySQL native function](#) in the *User Guide for Aurora*.

Oracle and MySQL character sets

With Amazon DMS, you can migrate databases between different platforms while preserving character set encodings and collations. Character sets define the encoding used to represent characters, while collations determine the sorting order and comparison rules. Properly configuring character sets and collations is crucial for applications that handle multilingual data or have specific sorting requirements.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	Different syntax. MySQL can have

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
			different collations for each database in the same instance.

Oracle usage

Oracle supports most national and international encoded character set standards including extensive support for Unicode.

Oracle provides two scalar string-specific data types:

- **VARCHAR2** — Stores variable-length character strings with a length between 1 and 4000 bytes. The Oracle database can be configured to use the VARCHAR2 data type to store either Unicode or Non-Unicode characters.
- **NVARCHAR2** — Scalar data type used to store Unicode data. Supports AL16UTF16 or UTF8 and is specified during database creation.

Character sets in Oracle are defined at the instance level (Oracle 11g) or the pluggable database level (Oracle 12c R2). In Pre-12cR2 Oracle databases, the character set for the root container and all pluggable databases were required to be identical.

Oracle 18c updates AL32UTF8 and AL16UTF16 character sets to Unicode standard version 9.0.

UTF8 Unicode

In Oracle, you can use the AL32UTF8 character set. Oracle provides encoding of ASCII characters as single-byte for Latin characters, two-bytes for some European and Middle-Eastern languages, and three-bytes for certain South and East-Asian characters. Therefore, Unicode storage requirements are usually higher when compared non-Unicode character sets.

Character set migration

Two options exist for modifying existing Instance-level or database-level character sets:

- Export or import from the source Instance/PDB to a new Instance/PDB with a modified character set.
- Use the Database Migration Assistant for Unicode (DMU), which simplifies the migration process to the Unicode character set.

As of 2012, use of the CSALTER utility for character set migrations is deprecated.

Oracle Database 12c Release 1 (12.1.0.1) complies with version 6.1 of the Unicode standard.

Oracle Database 12c Release 2 (12.1.0.2) extends the compliance to version 6.2 of the Unicode standard.

UTF-8 is supported through the AL32UTF8 CS and is valid as both the client and database character sets.

UTF-16BE is supported through AL16UTF16 and is valid as the national (NCHAR) character set.

For more information, see [Choosing a Character Set](#), [Locale Data](#), and [Supporting Multilingual Databases with Unicode](#) in the *Oracle documentation*.

MySQL usage

MySQL supports a variety of different character sets including support for both single-byte and multi-byte languages. The default character set is specified when initializing a MySQL database cluster with `initdb`. Each individual database created on the MySQL cluster supports individual character sets defined as part of database creation.

To query the available character sets, use the `INFORMATION_SCHEMA CHARACTER_SETS` table or the `SHOW CHARACTER SET` statement.

All character sets have at least one collation, and most character sets have more. To list the display collations for a character set, use the `INFORMATION_SCHEMA COLLATIONS` table or the `SHOW COLLATION` statement.

Collations have these general characteristics:

- Two different character sets cannot have the same collation.
- Each character set has a default collation.

- Collation names start with the name of the character set with which they are associated and are generally followed by one or more suffixes indicating other collation characteristics.

Examples

Create a database named test01 which uses the Korean EUC_KR Encoding the and the ko_KR locale.

```
CREATE DATABASE test01 CHARACTER SET = euckr COLLATE = euckr_korean_ci;
```

View the character sets configured for each database by querying the System Catalog.

```
SELECT SCHEMA_NAME,
       DEFAULT_CHARACTER_SET_NAME,
       DEFAULT_COLLATION_NAME
FROM INFORMATION_SCHEMA.SCHEMATA;
```

Convert a character set and collation using the ALTER DATABASE command.

```
ALTER DATABASE test01 CHARACTER SET = ucs2 COLLATE = ucs2_general_ci;
```

MySQL supports conversion of character sets between server and client for specific character set combinations with the parameter `character_set_client` and `character_set_connection`. For more information, see [Connection Character Sets and Collations](#).

In MySQL, you can specify the sort order and character classification behavior on a per-column level. Specify specific collations for individual table columns.

```
CREATE TABLE lang(
  latin1_col CHAR(10) CHARACTER SET latin1 COLLATE latin1_german1_ci,
  latin2_col CHAR(10) CHARACTER SET latin2);
```

Summary

Feature	Oracle	Aurora MySQL
View database character set	<pre>SELECT * FROM NLS_DATAB ASE_PARAMETERS;</pre>	<pre>SELECT SCHEMA_NAME,</pre>

Feature	Oracle	Aurora MySQL
		<pre> DEFAULT_CHARACTER_ SET_NAME, DEFAULT_COLLATION_ NAME FROM INFORMATION_SCHEMA .SCHEMATA; </pre>
Modify the database character set	<p>Choose one of the following options:</p> <ol style="list-style-type: none"> 1. Full export or import. 2. When converting to Unicode, use the Oracle DMU utility. 	<pre> ALTER DATABASE test01 CHARACTER SET = ucs2 COLLATE = ucs2_gene ral_ci; </pre>
Character set granularity	<p>Instance (11g + 12cR1)</p> <p>Database (Oracle 12cR2)</p>	Column
UTF8	Supported by using VARCHAR2 and NVARCHAR	Supported by using CHAR and VARCHAR
UTF16	Supported by using NVARCHAR2	Supported by using CHAR and VARCHAR
NCHAR and NVARCHAR data types	Supported	Supported

For more information, see [Character Sets, Collations, Unicode](#) and [Database Character Set and Collation](#) in the *MySQL documentation*.

Oracle database links and MySQL fully-qualified table names

With Amazon DMS, you can migrate data between different database platforms, including Oracle and MySQL, while preserving database links and fully-qualified table names. Oracle database links provide a way to access data in remote databases, while MySQL fully-qualified table names specify the database and table for a given object.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	MySQL doesn't support database links.

Oracle usage

Database links are schema objects used to interact with remote database objects such as tables. Common use cases for database links include selecting data from tables that reside in a remote database.

To use database links, Oracle net services must be installed on both the local and remote database servers to facilitate communications.

Examples

Create a database link named `remote_db`. When creating a database link, you have the option to specify the remote database destination using a TNS Entry or to specify the full TNS Connection string.

```
CREATE DATABASE LINK remote_db CONNECT TO username IDENTIFIED BY password USING
'remote';
```

```
CREATE DATABASE LINK remotenoTNS CONNECT TO username IDENTIFIED BY password
USING '(DESCRIPTION=(ADDRESS_LIST=(ADDRESS = (PROTOCOL = TCP)(HOST =192.168.1.1)
(PORT =1521)))(CONNECT_DATA =(SERVICE_NAME = orcl)))';
```

After the database link is created, you can use the database link directly as part of a SQL query using the database link name `@remote_db` as a suffix to the table name.

```
SELECT * FROM employees@remote_db;
```

Database links also support DML commands.

```
INSERT INTO employees@remote_db
```

```
(employee_id, last_name, email, hire_date, job_id) VALUES  
(999, 'Claus', 'sclaus@example.com', SYSDATE, 'SH_CLERK');  
  
UPDATE jobs@remote_db SET min_salary = 3000 WHERE job_id = 'SH_CLERK';  
  
DELETE FROM employees@remote_db WHERE employee_id = 999;
```

For more information, see [Managing Database Links](#) in the *Oracle documentation*.

MySQL usage

Currently, MySQL doesn't provide a direct comparable alternative for Oracle Database Links. You can use the fully-qualified names to query data from another database within the same cluster. This functionality is similar to querying data from a different schema in Oracle. If the data cannot be stored under the same MySQL Cluster, then there is no equivalent to Oracle Database Links in MySQL.

If the data can't be placed under the same MySQL Cluster then there is no relevant equivalent to Oracle Database Links in MySQL.

Examples

Query all flight ids from the `all_flights` table in the `flights` database, assume that this code runs from another database.

```
SELECT flight_id from flights.all_flights;
```

This query returns the data only if the user has permissions to the table and the database.

Oracle DBMS_SCHEDULER and MySQL events

With Amazon DMS, you can schedule and automate database tasks using Oracle DBMS_SCHEDULER and MySQL events. Oracle DBMS_SCHEDULER is an enterprise job scheduler that provides a way to schedule and automate recurring database tasks. MySQL events are similar, allowing you to schedule statements or stored procedures to execute at a specific time or interval.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Different paradigm and syntax

Oracle usage

The DBMS_SCHEDULER package contains a collection of scheduling functions that can be called from PL/SQL.

There are two main objects involved with creating scheduling jobs: a program and schedule. A program defines what to run, and a schedule defines when to run the program. The scheduler can run a database program unit such as a procedure or an external executable such as files system shell scripts.

There are three running methods for jobs: time-based scheduling, event-based jobs, and dependency jobs or chained jobs.

Time-based scheduling

The following examples create a job with a program and a schedule.

1. Create a program that will call the UPDATE_HR_SCHEMA_STATS procedure in the HR schema.
2. Create a schedule that will set the interval of running the jobs that using it. This schedule will run the job every hour.
3. Create the job.

```
BEGIN
DBMS_SCHEDULER.CREATE_PROGRAM(
program_name => 'CALC_STATS',
program_action => 'HR.UPDATE_HR_SCHEMA_STATS',
program_type => 'STORED_PROCEDURE',
enabled => TRUE);
END;
/
```

```
BEGIN
DBMS_SCHEDULER.CREATE_SCHEDULE(
  schedule_name => 'stats_schedule',
  start_date => SYSTIMESTAMP,
  repeat_interval => 'FREQ=HOURLY;INTERVAL=1',
  comments => 'Every hour');
END;
/

BEGIN
DBMS_SCHEDULER.CREATE_JOB (
  job_name => 'my_new_job3',
  program_name => 'my_saved_program1',
  schedule_name => 'my_saved_schedule1');
END;
/
```

Create a job without a program or a schedule:

- `job_type`: EXECUTABLE — The job runs as an external script.
- `job_action` — Defines the location of the external script.
- `start_date` — Defines when the job will be turned on.
- `repeat_interval` — Defines when the job will run. In the following example, the job runs every day at 23:00.

```
BEGIN
DBMS_SCHEDULER.CREATE_JOB(
  job_name=>'HR. BACKUP',
  job_type => 'EXECUTABLE',
  job_action => '/home/user/dba/rman/nightly_bck.sh',
  start_date=> SYSDATE,
  repeat_interval=>'FREQ=DAILY;BYHOUR=23',
  comments => 'Nightly backups');
END;
/
```

After you created the job, you can update its attributes with the `SET_ATTRIBUTE` procedure.

```
BEGIN
```

```
DBMS_SCHEDULER.SET_ATTRIBUTE (  
name => 'my_emp_job1',  
attribute => 'repeat_interval',  
value => 'FREQ=DAILY');  
END;  
/
```

Event-based jobs

The following example demonstrates how to create a schedule to start a job whenever the scheduler receives an event indicating a file arrived on the system before 9:00, and then create a job to use the schedule.

```
BEGIN  
DBMS_SCHEDULER.CREATE_EVENT_SCHEDULE (  
schedule_name => 'scott.file_arrival',  
start_date => systimestamp,  
event_condition => 'tab.user_data.object_owner = ''SCOTT''  
and tab.user_data.event_name = ''FILE_ARRIVAL''  
and extract hour from tab.user_data.event_timestamp < 9',  
queue_spec => 'my_events_q');  
END;  
/  
  
BEGIN  
DBMS_SCHEDULER.CREATE_JOB (  
job_name => my_job,  
program_name => my_program,  
start_date => '15-JUL-04 1.00.00AM US/Pacific',  
event_condition => 'tab.user_data.event_name = ''LOW_INVENTORY''',  
queue_spec => 'my_events_q'  
enabled => TRUE,  
comments => 'my event-based job');  
END;  
/
```

Dependency jobs

1. Use `DBMS_SCHEDULER.CREATE_CHAIN` to create a chain.
2. Use `DBMS_SCHEDULER.DEFINE_CHAIN_STEP`` to define three steps for this chain. Referenced programs must be enabled.

3. Use `DBMS_SCHEDULER.DEFINE_CHAIN_RULE` to define corresponding rules for the chain.
4. Use `DBMS_SCHEDULER.ENABLE` to enable the chain.
5. Use `DBMS_SCHEDULER.CREATE_JOB` to create a chain job to start the chain daily at 1:00 p.m.

```
BEGIN
DBMS_SCHEDULER.CREATE_CHAIN (
chain_name => 'my_chain1',
rule_set_name => NULL,
evaluation_interval => NULL,
comments => NULL);
END;
/

BEGIN
DBMS_SCHEDULER.DEFINE_CHAIN_STEP('my_chain1', 'stepA', 'my_program1');
DBMS_SCHEDULER.DEFINE_CHAIN_STEP('my_chain1', 'stepB', 'my_program2');
DBMS_SCHEDULER.DEFINE_CHAIN_STEP('my_chain1', 'stepC', 'my_program3');
END;
/

BEGIN
DBMS_SCHEDULER.DEFINE_CHAIN_RULE ('my_chain1', 'TRUE', 'START stepA');
DBMS_SCHEDULER.DEFINE_CHAIN_RULE (
'my_chain1', 'stepA COMPLETED', 'Start stepB, stepC');
DBMS_SCHEDULER.DEFINE_CHAIN_RULE (
'my_chain1', 'stepB COMPLETED AND stepC COMPLETED', 'END');
END;
/

BEGIN
DBMS_SCHEDULER.ENABLE('my_chain1');
END;
/

BEGIN
DBMS_SCHEDULER.CREATE_JOB (
job_name => 'chain_job_1',
job_type => 'CHAIN',
job_action => 'my_chain1',
repeat_interval => 'freq=daily;byhour=13;byminute=0;bysecond=0',
enabled => TRUE);
END;
```

/

There are two additional objects associated with jobs.

- **JOB CLASS** — When you have a number of jobs that has the same behavior and attributes, you may want to group them together into a bigger logical group called job class and you can give priority between job classes by allocating a high percentage of available resources.
- **WINDOW** — When you want to prioritize your jobs based on schedule, you can create a window of time that the jobs can run during this window, for example, during non-peak time or at the end of the month.

For more information, see [Scheduling Jobs with Oracle Scheduler](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL can use the EVENT objects to run scheduled events in the database. It can run a one-time event or a repeated event. In this case, it's called cycled. A repeated event is a time-based trigger that runs SQL, runs commands, or calls a procedure.

To use this feature, make sure that the `event_scheduler` parameter is set to ON. This isn't the default value.

If an EVENT terminates with errors, it is written to the error log. If there is a need to simulate the `dba_scheduler_job_log`, you can define the error log to use TABLE as the output.

For more information, see [Oracle Alert Log and MySQL Error Log](#).

Examples

Check that the event scheduler process is turned on.

```
select @@GLOBAL.event_scheduler
```

View all events.

```
select * from INFORMATION_SCHEMA.EVENTS;
```

Create a new event that runs a procedure every minute.

```
CREATE EVENT event_exec_myproc ON SCHEDULE EVERY 1 MINUTE
```

```
DO CALL simpleproc1(5);
```

Summary

Description	Oracle Scheduler	MySQL Events
<p>Create a job that runs as a stored procedure</p>	<pre> BEGIN DBMS_SCHEDULER.CREATE _PROGRAM(program_name => 'CALC_STATS', program_action => 'HR.UPDATE_HR_SCHE MA_STATS', program_type => 'STORED_PROCEDURE', enabled => TRUE); END; / BEGIN DBMS_SCHEDULER. CREATE_SCHEDULE(schedule_name => 'stats_schedule', start_date => SYSTIMESTAMP, repeat_interval => 'FREQQ=HOURLY;INTE RVAL=1', comments => 'Every hour'); END; / BEGIN DBMS_SCHEDULER .CREATE_JOB (job_name => 'my_new_j ob3', program_name => 'my_saved_program1', </pre>	<pre> CREATE EVENT stats_sch edule ON SCHEDULE EVERY 1 HOUR DO CALL HR.UPDATE _HR_SCHEMA_STATS(); </pre>

Description	Oracle Scheduler	MySQL Events
	<pre> schedule_name => 'my_saved_schedule 1'); END; / </pre>	
<p>Create a job that runs external executables</p>	<pre> BEGIN DBMS_SCHEDULER.CREATE _PROGRAM (program_name => 'oe.my_saved_progr am1', program_action => '/ usr/local/bin/date', program_type => 'EXECUTABLE', comments => 'My comments here'); END; / </pre>	<p>Use the following code to run an Amazon Lambda function:</p> <pre> CALL mysql.lambda_async ('arn:aws:lambda:us -west-2:1234567890 12:function:oe.my_ saved_program1', '{"input1":"value" }') </pre> <p>For more information, see Invoking a Lambda function from an Amazon Aurora MySQL DB cluster in the <i>User Guide for Aurora</i>.</p> <p>The <code>lambda_async</code> function runs a Lambda function and gets a JSON object for the input values.</p>

Description	Oracle Scheduler	MySQL Events
Create an event-based job	<pre> BEGIN DBMS_SCHEDULER.CREATE _EVENT_SCHEDULE (schedule_name => 'scott.file_arrival', start_date => systimestamp, event_condition => 'tab.user_data.obj ect_owner = ''SCOTT'' and tab.user_ data.event_name = ''FILE_ARRIVAL'' and extract hour from tab.user_ data.event_timestamp < 9', queue_spec => 'my_events_q'); END; / BEGIN DBMS_SCHEDULER. CREATE_JOB (job_name => my_job, program_name => my_program, start_date => '15- JUL-04 1.00.00AM US/ Pacific', event_condition => 'tab.user_data.eve nt_name = ''LOW_INV ENTORY''', queue_spec => 'my_events_q' enabled => TRUE, comments => 'my event- based job'); END; </pre>	<p>For the CREATE EVENT syntax, only time intervals can be defined as triggers for the event.</p> <p>If an event job is required, the best alternatives are:</p> <ol style="list-style-type: none"> 1. Create triggers to run the commands (for DML events). 2. Create an EVENT that runs every X time and check if the event occurred. The minimum interval is one second.

Description	Oracle Scheduler	MySQL Events
	/	

Description	Oracle Scheduler	MySQL Events
Create a chained job	<pre> BEGIN DBMS_SCHEDULER.CREATE _CHAIN (chain_name => 'my_chain1', rule_set_name => NULL, evaluation_interval => NULL, comments => NULL); END; / BEGIN DBMS_SCHEDULER. DEFINE_CHAIN_STEP ('my_chain1', 'stepA', 'my_program1'); DBMS_SCHEDULER.DEFIN E_CHAIN_STEP ('my_chain1', 'stepB', 'my_program2'); DBMS_SCHEDULER.DEFIN E_CHAIN_STEP ('my_chain1', 'stepC', 'my_program3'); END; / BEGIN DBMS_SCHEDULER. DEFINE_CHAIN_RULE ('my_chain1', 'TRUE', 'START stepA'); DBMS_SCHEDULER.DEF INE_CHAIN_RULE ('my_chain1', 'stepA COMPLETED', 'Start stepB, stepC'); DBMS_SCHEDULER.DEF INE_CHAIN_RULE (</pre>	Create several EVENTS and manage them within a table to keep the results, or the last run status to determine when to execute the next event.

Description	Oracle Scheduler	MySQL Events
	<pre> 'my_chain1', 'stepB COMPLETED AND stepC COMPLETED', 'END'); END; / BEGIN DBMS_SCHEDULER .ENABLE('my_chain1'); END; / BEGIN DBMS_SCHEDULER.CRE ATE_JOB (job_name => 'chain_jo b_1', job_type => 'CHAIN', job_action => 'my_chain1', repeat_interval => 'freq=daily; byhour=13; byminute=0; bysecond=0', enabled => TRUE); END; / </pre>	

For more information, see [Using the Event Scheduler](#) and [Event Syntax](#) in the *MySQL documentation*.

Oracle external tables and MySQL integration with Amazon S3

With Amazon DMS, you can integrate Oracle databases and MySQL databases with Amazon S3 using external tables and the Amazon-developed MySQL integration. External tables allow Oracle databases to directly access and query data stored in Amazon S3, while the MySQL integration lets MySQL databases treat Amazon S3 as a storage engine.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Creating Tables	Use Aurora MySQL integration with Amazon S3. Different paradigm and syntax.

Oracle usage

The Oracle external tables feature allows you to create a table in your database that reads data from a source located outside your database (externally).

Beginning with Oracle 12.2, the external table can be partitioned, providing all the benefits of a regular partitioned table.

Oracle 18c adds support for inline external tables, which is a way to get data from external source in a SQL query without having to define and create external table first.

```
SELECT * FROM EXTERNAL ((i NUMBER, d DATE)
TYPE ORACLE_LOADER
DEFAULT DIRECTORY data_dir
ACCESS PARAMETERS (
RECORDS DELIMITED BY NEWLINE
FIELDS TERMINATED BY '|' )
LOCATION ('test.csv')
REJECT LIMIT UNLIMITED)
tst_external;
```

Examples

Use `CREATE TABLE` with `ORGANIZATION EXTERNAL` to identify it as an external table. Specify the `TYPE` to let the database choose the right driver for the data source, the options are:

- `ORACLE_LOADER` — The data must be sourced from text data files. This is the default option.
- `ORACLE_DATAPUMP` — The data must be sourced from binary dump files. You can write dump files only as part of creating an external table with the `CREATE TABLE AS SELECT` statement.

Once the dump file is created, it can be read any number of times, but it can't be modified. This means that no DML operations can be performed.

- ORACLE_HDFS — Extracts data stored in a Hadoop Distributed File System (HDFS).
- ORACLE_HIVE — Extracts data stored in Apache HIVE.
- DEFAULT DIRECTORY — In database definition for the directory path.
- ACCESS PARAMETER — Defines the delimiter character and the query fields.
- LOCATION — The file name in the first two data source types or URI in the Hadoop data source (not in use with hive data source).

```
CREATE TABLE emp_load
(id CHAR(5), emp_dob CHAR(20), emp_lname CHAR(30),
 emp_fname CHAR(30), emp_start_date DATE) ORGANIZATION EXTERNAL
(TYPE ORACLE_LOADER DEFAULT DIRECTORY data_dir ACCESS PARAMETERS
(RECORDS DELIMITED BY NEWLINE FIELDS (id CHAR(2), emp_dob CHAR(20),
 emp_lname CHAR(18), emp_fname CHAR(11), emp_start_date CHAR(10)
 date_format DATE mask "mm/dd/yyyy")))
LOCATION ('info.dat'));
```

For more information, see [External Tables Concepts](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL has a capability similar to Oracle's External Tables, but requires a significant amount of syntax modifications. The main difference is that there is no open link to files and the data must be transferred from and to MySQL if you need all data.

There are two important operations for MySQL and S3 integration:

- Saving data to an S3 file.
- Loading data from an S3 file.

Aurora MySQL must have permissions to the S3 bucket.

In Oracle 18c, the inline external table feature was introduced. This cannot be achieved in Aurora for MySQL and it depends on the use case but other services can be considered. For ETLs, for example, Amazon Glue can be considered.

Saving data to Amazon S3

You can use the `SELECT INTO OUTFILE S3` statement to query data from an Amazon Aurora MySQL DB cluster and save it directly into text files stored in an Amazon S3 bucket. Use this functionality to avoid transferring data to the client first, and then copying the data from the client to Amazon S3.

Note

The default file size threshold is six gigabytes (GB). If the data selected by the statement is less than the file size threshold, a single file is created. Otherwise, multiple files are created.

If the `SELECT` statement fails, files that are already uploaded to Amazon S3 remain in the specified Amazon S3 bucket. You can use another statement to upload the remaining data instead of starting over again.

If the amount of data to be selected is more than 25 GB, it is recommended to use multiple `SELECT INTO OUTFILE S3` statements to save data to Amazon S3.

Metadata, such as table schema or file metadata, isn't uploaded by Aurora MySQL to Amazon S3.

Examples

The following statement selects all data in the `employees` table and saves the data into an Amazon S3 bucket in a different region from the Aurora MySQL DB cluster. The statement creates data files in which each field is terminated by a comma `,` character and each row is terminated by a newline `\n` character. The statement returns an error if files that match the `sample_employee_data` file prefix exist in the specified Amazon S3 bucket.

```
SELECT * FROM employees INTO OUTFILE S3
's3-us-west-2://aurora-select-into-s3-pdx/sample_employee_data'
FIELDS TERMINATED BY ',' LINES TERMINATED BY '\n';
```

The following statement selects all data in the `employees` table and saves the data into an Amazon S3 bucket in the same region as the Aurora MySQL DB cluster. The statement creates data files in which each field is terminated by a comma `,` character and each row is terminated by a newline `\n` character. It also creates a manifest file. The statement returns an error if files that match the `sample_employee_data` file prefix exist in the specified Amazon S3 bucket.

```
SELECT * FROM employees INTO OUTFILE S3
's3://aurora-select-into-s3-pdx/sample_employee_data'
FIELDS TERMINATED BY ',' LINES TERMINATED BY '\n'
MANIFEST ON;
```

The following statement selects all data in the employees table and saves the data into an Amazon S3 bucket in a different region from the Aurora DB cluster. The statement creates data files in which each field is terminated by a comma , character and each row is terminated by a newline \n character. The statement overwrites any existing files that match the sample_employee_data file prefix in the specified Amazon S3 bucket.

```
SELECT * FROM employees INTO OUTFILE S3
's3-us-west-2://aurora-select-into-s3-pdx/sample_employee_data'
FIELDS TERMINATED BY ',' LINES TERMINATED BY '\n' OVERWRITE ON;
```

The following statement selects all data in the employees table and saves the data into an Amazon S3 bucket in the same region as the Aurora MySQL DB cluster. The statement creates data files in which each field is terminated by a comma , character and each row is terminated by a newline \n character. It also creates a manifest file. The statement overwrites any existing files that match the sample_employee_data file prefix in the specified Amazon S3 bucket.

```
SELECT * FROM employees INTO OUTFILE S3
's3://aurora-select-into-s3-pdx/sample_employee_data'
FIELDS TERMINATED BY ',' LINES TERMINATED BY '\n'
MANIFEST ON OVERWRITE ON;
```

For more information, see [Saving data from an Amazon Aurora MySQL DB cluster into text files in an Amazon S3 bucket](#) in the *User Guide for Aurora*.

Loading data from Amazon S3

You can use the `LOAD DATA FROM S3` or `LOAD XML FROM S3` statement to load data from files stored in an Amazon S3 bucket.

Also, you can use the `LOAD DATA FROM S3` statement to load data from any text file format supported by the MySQL `LOAD DATA INFILE` statement such as comma-delimited text data. Compressed files aren't supported.

Examples

The following example runs the `LOAD DATA FROM S3` statement with the manifest from the previous example. This manifest has the `customer.manifest` name. After the statement completes, an entry for each successfully loaded file is written to the `aurora_s3_load_history` table.

```
LOAD DATA FROM S3 MANIFEST
's3-us-west-2://aurora-bucket/customer.manifest'
INTO TABLE CUSTOMER FIELDS TERMINATED BY ','
LINES TERMINATED BY '\n'
(ID, FIRSTNAME, LASTNAME, EMAIL);
```

Each successful `LOAD DATA FROM S3` statement updates the `aurora_s3_load_history` table in the `mysql` schema with an entry for each file loaded.

After you run the `LOAD DATA FROM S3` statement, you can verify which files were loaded by querying the `aurora_s3_load_history` table. To see the files that were loaded from one execution of the statement, use the `WHERE` clause to filter the records on the Amazon S3 URI for the manifest file used in the statement. If you have used the same manifest file before, filter the results using the timestamp field.

```
select * from mysql.aurora_s3_load_history
where load_prefix = 'S3_URI';
```

The following table describes the fields in the `aurora_s3_load_history` table:

Field	Description
<code>load_prefix</code>	<p>The URI specified in the load statement. This URI can map to any of the following:</p> <ul style="list-style-type: none"> A single data file for a <code>LOAD DATA FROM S3 FILE</code> statement. An Amazon S3 prefix that maps to multiple data files for a <code>LOAD DATA FROM S3 PREFIX</code> statement. A single manifest file containing the names of files to be loaded for a <code>LOAD DATA FROM S3 MANIFEST</code> statement.

Field	Description
file_name	The name of a file loaded into Aurora from Amazon S3 using the URI identified in the load_prefix field.
version_number	The version number of the file identified by the file_name field that was loaded if the Amazon S3 bucket has a version number.
bytes_loaded	The size of the file loaded in bytes.
load_timestamp	The timestamp when the LOAD DATA FROM S3 statement completed.

Examples

The following statement loads data from an Amazon S3 bucket in the same region as the Aurora DB cluster. The statement reads the comma-delimited data in the customerdata.txt file in the dbbucket Amazon S3 bucket and then loads the data into the store-schema.customer-table table.

```
LOAD DATA FROM S3 's3://dbbucket/customerdata.csv'  
INTO TABLE store-schema.customer-table  
FIELDS TERMINATED BY ','  
LINES TERMINATED BY '\n'  
(ID, FIRSTNAME, LASTNAME, ADDRESS, EMAIL, PHONE);
```

The following statement loads data from an Amazon S3 bucket in a different region from the Aurora DB cluster. The statement reads the comma-delimited data from all files matching the employee-data object prefix in the mydata Amazon S3 bucket in the us-west-2 region and then loads data into the employees table.

```
LOAD DATA FROM S3 PREFIX  
's3-us-west-2://my-data/employee_data'  
INTO TABLE employees  
FIELDS TERMINATED BY ','  
LINES TERMINATED BY '\n'  
(ID, FIRSTNAME, LASTNAME, EMAIL, SALARY);
```

The following statement loads data from the files specified in a JSON manifest file named `q1_sales.json` into the `sales` table.

```
LOAD DATA FROM S3 MANIFEST
's3-us-west-2://aurora-bucket/q1_sales.json'
INTO TABLE sales
FIELDS TERMINATED BY ','
LINES TERMINATED BY '\n'
(MONTH, STORE, GROSS, NET);
```

Loading XML FROM S3

You can use the `LOAD XML FROM S3` statement to load data from XML files stored on an Amazon S3 bucket in one of three different XML formats as shown following.

Column names are attributes of a `<row>` element. The attribute value identifies the contents of the table field.

```
<row column1="value1" column2="value2" .../>
```

Column names are child elements of a `<row>` element. The value of the child element identifies the contents of the table field.

```
<row>
<column1>value1</column1>
<column2>value2</column2>
</row>
```

Column names are in the name attribute of `<field>` elements in a `<row>` element. The value of the `<field>` element identifies the contents of the table field.

```
<row>
<field name='column1'>value1</field>
<field name='column2'>value2</field>
</row>
```

The following statement loads the first column from the input file into the first column of `table1` and sets the value of the `table_column2` column in `table1` to the input the value of the second column divided by 100.

```
LOAD XML FROM S3
's3://mybucket/data.xml'
INTO TABLE table1 (column1, @var1)
SET table_column2 = @var1/100;
```

The following statement sets the first two columns of `table1` to the values in the first two columns from the input file and then sets the value of the `column3` in `table1` to the current time stamp.

```
LOAD XML FROM S3
's3://mybucket/data.xml'
INTO TABLE table1 (column1, column2)
SET column3 = CURRENT_TIMESTAMP;
```

You can use subqueries in the right side of SET assignments. For a subquery that returns a value to be assigned to a column, you can use only a scalar subquery. Also, you can't use a subquery to select from the table that is being loaded.

For more information, see [Loading data into an Amazon Aurora MySQL DB cluster from text files in an Amazon S3 bucket](#) in the *Amazon RDS user guide*.

Oracle and MySQL inline views

With Amazon DMS, you can migrate data from Oracle and MySQL databases to Aurora or using inline views. Inline views are virtual tables that combine data from one or more tables using a SELECT statement. They are useful for querying complex data sets, joining data across tables, and simplifying data transformations during migration.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	N/A

Oracle usage

Inline views refer to SELECT statements located in the FROM clause of secondary SELECT statement. Inline views can help make complex queries simpler by removing compound calculations or eliminating join operations while condensing several separate queries into a single simplified query.

Examples

The SQL statement marked in red represents the inline view code. The query returns each employee matched to their salary and department id. In addition, the query returns the average salary for each department using the inline view column SAL_AVG.

```
SELECT A.LAST_NAME, A.SALARY, A.DEPARTMENT_ID, B.SAL_AVG
FROM EMPLOYEES A,
(SELECT DEPARTMENT_ID, ROUND(AVG(SALARY))
AS SAL_AVG FROM EMPLOYEES GROUP BY DEPARTMENT_ID)
WHERE A.DEPARTMENT_ID = B.DEPARTMENT_ID;
```

MySQL usage

MySQL semantics may refer to inline views as sub select or as subquery. In either case, the functionality is the same. Running the preceding Oracle inline view example, as is, will result in the following error: SQL Error[1248][4200]: Every derived table must have its own alias. This error occurs because Oracle supports omission of aliases for the inner statement while in MySQL aliases are mandatory. Mandatory aliases are the only major difference when migrating Oracle inline views to MySQL.

Examples

The following example uses B as an alias.

```
SELECT A.LAST_NAME, A.SALARY, A.DEPARTMENT_ID, B.SAL_AVG
FROM EMPLOYEES A,
(SELECT DEPARTMENT_ID, ROUND(AVG(SALARY)) AS SAL_AVG
FROM EMPLOYEES GROUP BY DEPARTMENT_ID) B
WHERE A.DEPARTMENT_ID = B.DEPARTMENT_ID;
```

Oracle JSON document support and MySQL JSON

With Amazon DMS, you can migrate data between different database platforms, including Oracle and MySQL, while preserving the JSON document structure. Oracle JSON document support and MySQL JSON provide a way to store and query JSON data within the database.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	Different paradigm and syntax will require application or drivers rewrite.

Oracle usage

JSON documents are based on JavaScript syntax and allow the serialization of objects. Oracle support for JSON document storage and retrieval enables you to extend the database capabilities beyond purely relational use cases and allows an Oracle database to support semi-structured data. Oracle JSON support also includes full-text search and several other functions dedicated to querying JSON documents.

Oracle 19 adds a new `JSON_SERIALIZE` function. You can use this function to serialize JSON objects to text.

For more information, see [Introduction to JSON Data and Oracle Database](#) in the *Oracle documentation*.

Examples

The following example creates a table to store a JSON document in a data column and insert a JSON document into the table.

```
CREATE TABLE json_docs (id RAW(16) NOT NULL, data CLOB,
CONSTRAINT json_docs_pk PRIMARY KEY (id),
CONSTRAINT json_docs_json_chk CHECK (data IS JSON));
```

```
INSERT INTO json_docs (id, data) VALUES (SYS_GUID(),
'{'
  "FName" : "John",
  "LName" : "Doe",
  "Address" : {
    "Street" : "101 Street",
    "City" : "City Name",
    "Country" : "US",
    "Pcode" : "90210"}
}');
```

Unlike XML data, which is stored using the XMLType SQL data type, JSON data is stored in an Oracle Database using the SQL data types VARCHAR2, CLOB, and BLOB. Oracle recommends that you always use an `is_json` check constraint to ensure the column values are valid JSON instances. Or, add a constraint at the table-level `CONSTRAINT json_docs_json_chk CHECK (data IS JSON)`.

You can query a JSON document directly from a SQL query without the use of special functions. Querying without functions is called Dot Notation.

```
SELECT a.data.FName,a.data.LName,a.data.Address.Pcode AS Postcode
FROM json_docs a;
```

FNAME	LNAME	POSTCODE
John	Doe	90210

1 row selected.

In addition, Oracle provides multiple SQL functions that integrate with the SQL language and enable querying JSON documents (such as `IS JSON`, `JSON_VALUE`, `JSON_EXISTS`, `JSON_QUERY`, and `JSON_TABLE`).

For more information, see [Introduction to JSON Data and Oracle Database](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL 5.7 supports a native JSON data type for storing JSON documents, which provides several benefits over storing the same document as a generic string. All JSON documents stored as a JSON data type are validated for correctness. If the document is not valid JSON, it is rejected and an error condition is raised. In addition, more efficient storage algorithms enable optimized

read access to elements within the document. The optimized internal binary representation of the document enables much faster operation on the data without requiring expensive re-parsing.

Consider the following example:

```
CREATE TABLE JSONTable (  
  DocumentIdentifier INT NOT NULL PRIMARY KEY,  
  JSONDocument JSON);
```

MySQL 5.7.22 also added the JSON utility function `JSON_PRETTY()` which outputs an existing JSON value in an easy-to-read format; each JSON object member or array value is printed on a separate line and a child object or array is indented two spaces with respect to its parent. This function also works with a string that can be parsed as a JSON value. For more information, see [JSON Utility Functions](#) in the *MySQL documentation*.

MySQL 5.7.22 also added the JSON utility functions `JSON_STORAGE_SIZE()` and `JSON_STORAGE_FREE()`.

`JSON_STORAGE_SIZE()` returns the storage space in bytes used for the binary representation of a JSON document prior to any partial update.

`JSON_STORAGE_FREE()` shows the amount of space freed after it has been partially updated using `JSON_SET()` or `JSON_REPLACE()`. This is greater than zero if the binary representation of the new value is less than that of the previous value. Each of these functions also accepts a valid string representation of a JSON document. For such a value `JSON_STORAGE_SIZE()` returns the space used by its binary representation following its conversion to a JSON document. For a variable containing the string representation of a JSON document `JSON_STORAGE_FREE()` returns zero.

These functions produce an error if the non-null argument can't be parsed as a valid JSON document and NULL if the argument is NULL. For more information, see [JSON Utility Functions](#) in the *MySQL documentation*.

Note

Amazon Relational Database Service (Amazon RDS) for MySQL version 8 added two JSON aggregation functions `JSON_ARRAYAGG()` and `JSON_OBJECTAGG()`. `JSON_ARRAYAGG()` takes a column or expression as its argument and aggregates the result as a single JSON array. The expression can evaluate to any MySQL data type; this does not have to be a JSON value.

`JSON_OBJECTAGG()` takes two columns or expressions which it interprets as a key and a value; it returns the result as a single JSON object.

Note

Amazon RDS for MySQL version 8.0.17 adds two functions `JSON_SCHEMA_VALID()` and `JSON_SCHEMA_VALIDATION_REPORT()` for validating JSON documents. `JSON_SCHEMA_VALID()` returns `TRUE` (1) if the document validates against the schema and `FALSE` (0) if it doesn't. `JSON_SCHEMA_VALIDATION_REPORT()` returns a JSON document containing detailed information about the results of the validation.

JSON functions

Aurora MySQL supports a rich set of more than 25 targeted functions for working with JSON data. These functions enable adding, modifying, and searching JSON data. Additionally, spatial JSON functions can be used for GeoJSON documents. For more information, see [Spatial GeoJSON Functions](#) in the *MySQL documentation*.

The `JSON_ARRAY`, `JSON_OBJECT`, and `JSON_QUOTE` functions all return a JSON document from a list of values, a list of key-value pairs, or a JSON value respectively.

Consider the following example:

```
SELECT JSON_OBJECT('Person', 'John', 'Country', 'USA');
{"Person": "John", "Country": "USA"}
```

The `JSON_CONTAINS`, `JSON_CONTAINS_PATH`, `JSON_EXTRACT`, `JSON_KEYS`, and `JSON_SEARCH` functions are used to query and search the content of a JSON document.

The `CONTAINS` functions are Boolean functions that return 1 or 0 (`TRUE` or `FALSE`).

`JSON_EXTRACT` returns a subset of the document based on the `XPATH` expression.

`JSON_KEYS` returns a JSON array consisting of the top-level key (or path top level) values of a JSON document.

The `JSON_SEARCH` function returns the path to one or all of the instances of the search string.

Examples

```
SELECT JSON_EXTRACT('["Mary", "Paul", ["Jim", "Ryan"]]', '$[1]');  
  
"Paul"  
  
SELECT JSON_SEARCH('["Mary", "Paul", ["Jim", "Ryan"]]', 'one', 'Paul');  
  
"$[1]"
```

Aurora MySQL supports the following functions for adding, deleting, and modifying JSON data: `JSON_INSERT`, `JSON_REMOVE`, `JSON_REPLACE`, and their `ARRAY` counterparts, which are used to create, delete, and replace existing data elements.

```
SELECT JSON_ARRAY_INSERT('["Mary", "Paul", "Jim"]', '$[1]', 'Jack');  
  
["Mary", "Jack", "Paul", "Jim"]
```

`JSON_SEARCH` is used to find the location of an element value within a JSON document.

```
SELECT JSON_SEARCH('["Mary", "Paul", ["Jim", "Ryan"]]', 'one', 'Paul');  
  
"$[1]"
```

JSON indexes

JSON columns are effectively a `BINARY` family type, which cannot be indexed. As an alternative, use `CREATE TABLE` or `ALTER TABLE` to add generated columns that represent some value from the JSON document and create an index on the generated column. For more information, see [Oracle Virtual Columns and MySQL Generated Columns](#).

Note

If indexes on generated columns exist for JSON documents, the query optimizer can use them to match JSON expressions and optimize data access.

Summary

Feature	Oracle	Aurora MySQL
JSON functions	IS_JSON, IS_NOT_JSON , JSON_EXISTS , JSON_VALUE , JSON_QUERY , JSON_TABLE	A set of more than 25 dedicated JSON functions . For more information, see JSON Function Reference in the <i>MySQL documentation</i> .
Return the full JSON document or all JSON documents	The emp_data column stores json documents: <pre>SELECT emp_data FROM employees;</pre>	The emp_data column stores json documents: <pre>SELECT emp_data FROM employees;</pre>
Return a specific element from a JSON document	Return only the address property: <pre>SELECT e.emp_data->'address' FROM employees e;</pre>	Return only the address property: <pre>SELECT emp_data->>'address' from employees where emp_id = 1;</pre>
Return JSON documents matching a pattern in any field	Return the JSON based on a search of on all JSON properties. Could be returned even if element is equal to the pattern. <pre>SELECT e.emp_data FROM employees e WHERE e.emp_data like '%pattern%';</pre>	Return the JSON based on a search of on all JSON properties. Could be returned even if element is equal to the pattern. <pre>SELECT e.emp_data FROM employees e WHERE e.emp_data like '%pattern%';</pre>
Return JSON documents matching a pattern in specific fields (root level)	<pre>SELECT e.emp_data.name FROM employees e</pre>	<pre>SELECT emp_data.name FROM employees</pre>

Feature	Oracle	Aurora MySQL
	<pre>WHERE e.data.active = 'true';</pre>	<pre>WHERE emp_data->>"\$.active" = 'true';</pre>

For more information, see [The JSON Data Type](#) and [JSON Functions](#) in the *MySQL documentation*.

Oracle materialized views and MySQL summary tables or views

With Amazon DMS, you can create Oracle materialized views and MySQL summary tables or views to improve query performance and data availability. Materialized views and summary tables are database objects that store pre-computed results of queries, reducing the need for complex calculations at runtime.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Materialized Views	MySQL doesn't support materialized views.

Oracle usage

Oracle materialized views are table segments where the contents are periodically refreshed based on the results of a stored query. Oracle materialized views are defined with specific queries and can be manually or automatically refreshed based on specific configurations. A materialized view runs its associated query and stores the results as a table segment.

Oracle materialized views are especially useful for:

- Replication of data across multiple databases.
- Data warehouse use cases.
- Increasing performance by persistently storing the results of complex queries as database tables.

Such as ordinary views, you can create materialized views with a SELECT query. The FROM clause of a materialized view query can reference tables, views, and other materialized views. The source objects that a materialized view uses as data sources are also called master tables (replication terminology) or detail tables (data warehouse terminology).

Immediate or deferred refresh

When you create materialized views, use the BUILD IMMEDIATE option can to instruct Oracle to immediately update the contents of the materialized view by running the underlying query. This is different from a deferred update where the materialized view is populated only on the first requested refresh.

Fast and complete refresh

You can use one of the two following options to refresh data in your materialized view.

- REFRESH FAST — Incremental data refresh. Only updates rows that have changed since the last refresh of the Materialized View instead of performing a complete refresh. This type of refresh fails if materialized view logs have not been created.
- COMPLETE — The table segment used by the materialized view is truncated (data is cleared) and repopulated by running the associated query.

Materialized view logs

When you create materialized views, use a materialized view log to instruct Oracle to store any changes performed by DML commands on the master tables that are used to refresh the materialized view, which provides faster materialized view refreshes.

Without materialized view logs, Oracle must re-run the query associated with the materialized view each time. This process is also known as a complete refresh. This process is slower compared to using materialized view logs.

Materialized view refresh strategy

You can use one of the two following strategies to refresh data in your materialized view.

- ON COMMIT — Refreshes the materialized view upon any commit made on the underlying associated tables.

- **ON DEMAND** — The refresh is initiated by a scheduled task or manually by the user.

Examples

The following example creates a simple Materialized View named `mv1` that runs a simple `SELECT` statement on the `employees` table.

```
CREATE MATERIALIZED VIEW mv1 AS SELECT * FROM hr.employees;
```

The following example creates a more complex materialized view using a database link (`remote`) to obtain data from a table located in a remote database. This materialized view also contains a subquery. The `FOR UPDATE` clause allows the materialized view to be updated.

```
CREATE MATERIALIZED VIEW foreign_customers FOR
UPDATE AS SELECT * FROM sh.customers@remote cu WHERE EXISTS
(SELECT * FROM sh.countries@remote co WHERE co.country_id = cu.country_id);
```

The following example creates a materialized view on two source tables: `times` and `products`. This approach enables `FAST` refresh of the materialized view instead of the slower `COMPLETE` refresh. Also, create a new materialized view named `sales_mv` which is refreshed incrementally `REFRESH FAST` each time changes in data are detected (`ON COMMIT`) on one or more of the tables associated with the materialized view query.

```
CREATE MATERIALIZED VIEW LOG ON times
WITH ROWID, SEQUENCE (time_id, calendar_year)
INCLUDING NEW VALUES;

CREATE MATERIALIZED VIEW LOG ON products
WITH ROWID, SEQUENCE (prod_id)
INCLUDING NEW VALUES;

CREATE MATERIALIZED VIEW sales_mv
BUILD IMMEDIATE
REFRESH FAST ON COMMIT
AS SELECT t.calendar_year, p.prod_id,
SUM(s.amount_sold) AS sum_sales
FROM times t, products p, sales s
WHERE t.time_id = s.time_id AND p.prod_id = s.prod_id
GROUP BY t.calendar_year, p.prod_id;
```

For more information, see [Basic Materialized Views](#) in the *Oracle documentation*.

MySQL usage

Oracle materialized views have no equivalent feature in MySQL, but other features can be used separately or combined to achieve similar functionality.

Make sure that you evaluate each case on its own merits, but options include:

- **Summary tables** — If your materialized view has many calculations and data manipulations, you can keep the results in tables and query the data without running all calculations on-the-fly. The data for these tables can be copied using triggers or events objects.
- **Views** — Aurora MySQL has a new Parallel Query mechanism that offloads some of the query operations to the storage level. This approach can greatly improve performance. In some cases, regular views can be used and may decrease some administration tasks. To evaluate this option, measure the performance and execution time of your SQL.

For more information, see [CREATE TABLE Statement](#), [Trigger Syntax and Examples](#), and [CREATE VIEW Statement](#) in the *MySQL documentation*.


Oracle multitenant and MySQL databases

With Amazon DMS, you can migrate Oracle multitenant databases and MySQL databases to Amazon Aurora.


Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Distribute load, applications, and users across multiple instances.

Oracle usage

Oracle 12c introduces a new multitenant architecture that provides the ability to create additional independent pluggable databases under a single Oracle instance. Prior to Oracle 12c, a single Oracle database instance only supported running a single Oracle database as shown in the following diagram.



Oracle 12c introduces a new multitenant container database (CDB) that supports one or more pluggable databases (PDB). The CDB can be thought of as a single superset database with multiple pluggable databases. The relationship between an Oracle instance and databases is now 1:N.



Oracle 18c adds following multitenant related features:

- **DBCA PDB Clone** — UI interface which allows cloning multiple pluggable databases (PDB).
- **Refreshable PDB Switchover** — An ability to switch roles between pluggable database clone and its original primary.

- **CDB Fleet Management** — An ability to group multiple container databases (CDB) into fleets that can be managed as a single logical database.


Oracle 19 introduced support to having more than one pluggable database (PDB) in a container database (CDB) in sharded environments.

Advantages of the Oracle 12c multitenant architecture

- You can use PDBs to isolate applications from one another.
- You can use PDBs as portable collection of schemas.
- You can clone PDBs and transport them to different CDBs/Oracle instances.
- Management of many databases (individual PDBs) as a whole.
- Separate security, users, permissions, and resource management for each PDB provides greater application isolation.
- Enables a consolidated database model of many individual applications sharing a single Oracle server.
- Provides an easier way to patch and upgrade individual clients and/or applications using PDBs.
- Backups are supported at both a multitenant container-level as well as at an individual PDB-level (both for physical and logical backups).

The Oracle multitenant architecture

- A multitenant CDB can support one or more PDBs.
- Each PDB contains its own copy of SYSTEM and application tablespaces.
- The PDBs share the Oracle Instance memory and background processes. The use of PDBs enables consolidation of many databases and applications into individual containers under the same Oracle instance.
- A single Root Container (CDB\$ROOT) exists in a CDB and contains the Oracle Instance Redo Logs, undo tablespace (unless Oracle 12.2 local undo mode is enabled), and control files.
- A single Seed PDB exists in a CDB and is used as a template for creating new PDBs.



CDB and PDB semantics

Container databases (CDB)

- Created as part of the Oracle 12c software installation.
- Contains the Oracle control files, its own set of system tablespaces, the instance undo tablespaces (unless Oracle 12.2 local undo mode is enabled), and the instance redo logs.
- Holds the data dictionary for the root container and for all of the PDBs.

Pluggable databases (PDB)

- An independent database that exists under a CDB. Also known as a container.
- Used to store application-specific data.
- You can create a pluggable database from a the `pdb$seed` (template database) or as a clone of an existing PDB.
- Stores metadata information specific to its own objects (data-dictionary).
- Has its own set of application data files, system data files, and tablespaces along with temporary files to manage objects.

Examples

List existing PDBs created in an Oracle CDB instance.

```
SHOW PDBS;
```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO

```
3          PDB1          READ WRITE  NO
```

Provision a new PDB from the template seed\$pdb.

```
CREATE PLUGGABLE DATABASE PDB2 admin USER ora_admin
IDENTIFIED BY ora_admin FILE_NAME_CONVERT=('/pdbseed/', '/pdb2/');
```

Alter a specific PDB to the READ/WRITE mode and verify the change.

```
ALTER PLUGGABLE DATABASE PDB2 OPEN READ WRITE;
```

```
SHOW PDBS;
```

CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	PDB1	READ WRITE	NO
4	PDB2	READ WRITE	NO

Clone a PDB from an existing PDB.

```
CREATE PLUGGABLE DATABASE PDB3
FROM PDB2 FILE_NAME_CONVERT= ('/pdb2/', '/pdb3/');
```

```
SHOW PDBS;
```


CON_ID	CON_NAME	OPEN MODE	RESTRICTED
2	PDB\$SEED	READ ONLY	NO
3	PDB1	READ WRITE	NO
4	PDB2	READ WRITE	NO
5	PDB3	MOUNTED	

For more information, see [Oracle Multitenant](#) in the *Oracle documentation*.

MySQL usage


Amazon Aurora MySQL offers a different and simplified architecture to manage and create a multitenant database environment. You can use Aurora MySQL to provide levels of functionality similar but not identical to those offered by Oracle PDBs by creating multiple databases under the same Aurora MySQL cluster and / or using separate Aurora clusters if total isolation of workloads is required.

You can create multiple MySQL databases under a single Amazon Aurora MySQL cluster.



Each Amazon Aurora cluster contains a primary instance that can accept both reads and writes for all cluster databases.

You can create up to 15 read-only nodes providing scale-out functionality for application reads and high availability.



An Oracle CDB/Instance is a high-level equivalent to an Amazon Aurora cluster, and an Oracle Pluggable Database (PDB) is equivalent to a MySQL database created inside the Amazon Aurora cluster. Not all features are comparable between Oracle 12c PDBs and Amazon Aurora.

Starting with Oracle 18c and 19c, you can use this feature for the following:

- PDB Clone
- Refreshable PDB Switchover
- CDB Fleet Management
- More than one pluggable database (PDB) in a container database (CDB) in sharded environments.

In the Amazon Cloud, these features can be achieved in many ways and each can be optimized using different services.

Cloning databases inside the MySQL instance is not so easy. For the same instance, you can use export and import.

To achieve similar functionality to Refreshable PDB Switchover, it depends on the use case but there are multiple options mostly depended on the required granularity:

- Databases in the same instance — you can do the failover using `CREATE DATABASE` statement when size and required downtime allow that and use an application failover to point to any of the databases.
- Database links and replication method — database links or Amazon DMS can be used to make sure there are two databases in two different instances that are in sync and have application failover to point to the other database when needed.

Managing CDB is actually very similar to the Amazon orchestration, as you can manage multiple Amazon RDS instances there (CDB) and databases inside (PDB), all monitored centrally and can be managed through the Amazon console or Amazon CLI.

Examples

Create a new database in MySQL using the `CREATE DATABASE` statement.

```
CREATE DATABASE db1;
CREATE DATABASE db2;
CREATE DATABASE db3;
```

List all databases created under an Amazon Aurora MySQL cluster.

```
SHOW DATABASES;

Database
information_schema
mysql
performance_schema
db1
db2
db3
sys
tmp
```

Independent database backups

Oracle 12c provides the ability to perform both logical backups using DataPump and physical backups using RMAN at both the CDB and PDB levels. Similarly, Aurora MySQL provides the ability

to perform logical backups on all or a specific database using mysqldump. However, for physical backups when using snapshots, the entire cluster and all databases are included in the snapshot. Backing up a specific database within the cluster is not supported.

This is usually not a concern because volume snapshots are extremely fast operations that occur at the storage infrastructure layer, incur minimal overhead, and operate at extremely fast speeds. However, the process of restoring a single MySQL database from an Aurora snapshot requires additional steps such as exporting the specific database after a snapshot restore and importing it back to the original Aurora cluster.

For more information, see [CREATE DATABASE Statement](#) in the *MySQL documentation*.

Oracle Resource Manager and dedicated Amazon Aurora MySQL clusters

With Amazon DMS, you can migrate Oracle Resource Manager databases to dedicated Aurora MySQL clusters.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Distribute load, applications, or users across multiple instances.

Oracle usage

Oracle Resource Manager enables enhanced management of multiple concurrent workloads running under a single Oracle database. Using Oracle Resource Manager, you can partition server resources for different workloads.

Resource Manager helps with sharing server and database resources without causing excessive resource contention and helps to eliminate scenarios involving inappropriate allocation of resources across different database sessions.

Oracle Resource Manager enables you to:

- Guarantee a minimum amount of CPU cycles for certain sessions regardless of other running operations.
- Distribute available CPU by allocating percentages of CPU time to different session groups.
- Limit the degree of parallelism of any operation performed by members of a user group.
- Manage the order of parallel statements in the parallel statement queue.
- Limit the number of parallel running servers that a user group can use.
- Create an active session pool. An active session pool consists of a specified maximum number of user sessions allowed to be concurrently active within a user group.
- Monitor used database/server resources by dictionary views.
- Manage runaway sessions or calls and prevent them from overloading the database.
- Prevent the running of operations that the optimizer estimates will run for a longer time than a specified limit.
- Limit the amount of time that a session can be connected but idle, thus forcing inactive sessions to disconnect and potentially freeing memory resources.
- Allow a database to use different resource plans, based on changing workload requirements.
- Manage CPU allocation when there is more than one instance on a server in an Oracle Real Application Cluster environment (also called instance caging).

Oracle Resource Manager introduces three concepts:

- **Consumer Group** — A collection of sessions grouped together based on resource requirements. The Oracle Resource Manager allocates server resources to resource consumer groups, not to the individual sessions.
- **Resource Plan** — Specifies how the database allocates its resources to different Consumer Groups. You will need to specify how the database allocates resources by activating a specific resource plan.
- **Resource Plan Directive** — Associates a resource consumer group with a plan and specifies how resources are to be allocated to that resource consumer group.

 **Note**

Only one Resource Plan can be active at any given time.
Resource Directives control the resources allocated to a Consumer Group belong to a Resource Plan.

The Resource Plan can refer to Subplans to create even more complex Resource Plans.

Examples

Create a simple Resource Plan. To use the Oracle Resource Manager, you need to assign a plan name to the `RESOURCE_MANAGER_PLAN` parameter. Using an empty string will disable the Resource Manager.

```
ALTER SYSTEM SET RESOURCE_MANAGER_PLAN = 'mydb_plan';
ALTER SYSTEM SET RESOURCE_MANAGER_PLAN = '';
```

You can create complex Resource Plans. A complex Resource Plan is one that is not created with the `CREATE_SIMPLE_PLAN` PL/SQL procedure and provides more flexibility and granularity.

```
BEGIN
DBMS_RESOURCE_MANAGER.CREATE_PLAN_DIRECTIVE (
PLAN => 'DAYTIME',
GROUP_OR_SUBPLAN => 'OLTP',
COMMENT => 'OLTP group',
MGMT_P1 => 75);
END;
/
```

For more information, see [Managing Resources with Oracle Database Resource Manager](#) in the *Oracle documentation*.

MySQL usage

MySQL doesn't have built-in resource management capabilities that are equivalent to the functionality provided by Oracle Resource Manager. However, due to the elasticity and flexibility provided by cloud economics, workarounds could be applicable and such capabilities might not be as of similar importance to monolithic on-premises databases.

The Oracle Resource Manager primarily exists because traditionally, Oracle databases were installed on very powerful monolithic servers that powered multiple applications simultaneously. The monolithic model made the most sense in an environment where the licensing for the Oracle database was per-CPU and where Oracle databases were deployed on physical hardware. In these scenarios, it made sense to consolidate as many workloads as possible into few servers. In cloud


databases, the strict requirement to maximize the usage of each individual server is often not as important and a different approach can be employed:

Individual Amazon Aurora clusters can be deployed, with varying sizes, each dedicated to a specific application or workload. Additional read-only Aurora Replica servers can be used to offload any reporting-style workloads from the master instance.

The following diagram shows the traditional Oracle model where maximizing the usage of each physical Oracle server was essential due to physical hardware constraints and the per-CPU core licensing model.



With Amazon Aurora, you can deploy separate and dedicated database clusters. Each cluster is dedicated to a specific application or workload creating isolation between multiple connected sessions and applications. The following diagram shows this architecture.



Each Amazon Aurora instance (primary or replica) can be scaled independently in terms of CPU and memory resources using the different instance types. Because multiple Amazon Aurora instances can be instantly deployed and much less overhead is associated with the deployment and management of Aurora instances when compared to physical servers, separating different workloads to different instance classes could be a suitable solution for controlling resource management.

For instance types and resources, see [Amazon EC2 Instance Types](#).

In addition, each Amazon Aurora primary or replica instance can also be directly accessed from your applications using its own endpoint. This capability is especially useful if you have multiple Aurora read-replicas for a given cluster and you wish to utilize different Aurora replicas to segment your workload.

Examples

Suppose that you were using a single Oracle Database for multiple separate applications and used Oracle Resource Manager to enforce a workload separation, allocating a specific amount of

server resources for each application. With Amazon Aurora, you might want to create multiple separate databases for each individual application. Adding additional replica instances to an existing Amazon Aurora cluster is easy.

1. Sign in to your Amazon console and choose **RDS**.
2. Choose **Databases** and select the Amazon Aurora cluster that you want to scale-out by adding an additional reader.
3. Choose **Actions** and then choose **Add reader**.
4. Select the instance class depending on the amount of compute resources your application requires.
5. Choose **Create Aurora Replica**.

Summary

Oracle Resource Manager	Amazon Aurora instances
Set the maximum CPU usage for a resource group	Create a dedicated Aurora Instance for a specific application
Limit the degree of parallelism for specific queries	N/A
Limit parallel runs	N/A
Limit the number of active sessions	Manually detect the number of connections that are open from a specific application and restrict connectivity either with database procedures or within the application Data Access Layer (DAL). <pre>select count(*) from information_schema.processlist where user='USER_NAME' and COMMAND<>'Sleep';</pre>
Restrict maximum runtime of queries	<pre>SET max_execution_time TO X;</pre>

Oracle Resource Manager	Amazon Aurora instances
Limit the maximum idle time for sessions	<p>Manually detect the number of connections that are open from a specific application and restrict connectivity either with database procedures or within the application DAL.</p> <pre data-bbox="829 443 1507 722">select count(*) from information_schema.processlist where user='USER_NAME' and COMMAND='Sleep' and TIME > X;</pre>
Limit the time that an idle session holding open locks can block other sessions	<p>Manually detect the number of connections that are open from a specific application and restrict connectivity either with database procedures or within the application DAL.</p> <pre data-bbox="829 974 1507 1213">select count(*) from information_schema.processlist where user='USER_NAME' and COMMAND='Sleep';</pre>
Use instance caging in a multi-node Oracle RAC Environment	<p>You can achieve similar capabilities by separating different applications to different Aurora clusters or, for read-only workloads, separate Aurora read replicas within the same Aurora cluster.</p>

Oracle SecureFile LOBs and MySQL large objects

With Amazon DMS, you can migrate data from Oracle and MySQL databases to other database engines, including large object data types like Oracle SecureFile LOBs and MySQL large objects. Oracle SecureFile LOBs and MySQL large objects are data types that store large amounts of unstructured data, such as text, images, audio, and video files.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	MySQL doesn't support SecureFiles, automation and compatibility refer only to LOBs.

Oracle usage

Large objects (LOB) is a mechanism for storing binary data in a database. Oracle 11g introduced SecureFile LOBs that provide more efficient storage. They are created using the `SECUREFILE` keyword as part of the `CREATE TABLE` statement.

The Primary benefits of using `SECUREFILE` lobs include:

- **Compression** — Uses Oracle advanced compression to analyze SecureFiles LOB data to save disk space.
- **De-Duplication** — Automatically detects duplicate LOB data within a LOB column or partition and reduces storage space by removing duplicates of repeating binary data.
- **Encryption** — Combined with Transparent Data Encryption (TDE).

Examples

The following example creates a table using a SecureFiles LOB column.

```
CREATE TABLE sf_tab (COL1 NUMBER, COL2_CLOB CLOB) LOB(COL2_CLOB)
STORE AS SECUREFILE;
```

The following example provides additional options for LOB compression during table creation.

```
CREATE TABLE sf_tab (COL1 NUMBER,COL2_CLOB CLOB) LOB(COL2_CLOB)
STORE AS SECUREFILE COMPRESS_LOB(COMPRESS HIGH);
```

For more information, see [Introduction to Large Objects and SecureFiles](#) in the *Oracle documentation*.

MySQL usage

MySQL doesn't support the advanced storage, security, and encryption options of Oracle SecureFile LOBs. MySQL supports regular LOB datatypes and provides stream-style access.

The four Binary Large Object (BLOB) types are: TINYBLOB, BLOB, MEDIUMBLOB, and LONGBLOB.

These types differ only in the maximum length of the values they can hold.

The four TEXT types are: TINYTEXT, TEXT, MEDIUMTEXT, and LONGTEXT.

BLOB values are treated as binary or byte strings. They have the binary character set, collation, and comparison. Sorting is based on the numeric values of the bytes in column values.

TEXT values are treated as non-binary or character strings. They have a character set other than binary. Values are sorted and compared based on the collation of the character set.

For TEXT columns, index entries are space-padded at the end. If the index requires unique values, duplicate-key errors occur for values that differ only in the number of trailing spaces. For example, if a table contains 'b', an attempt to store 'b ' causes a duplicate-key error.

Because BLOB and TEXT values can be extremely long, there are some constraints:

- Only the first `max_sort_length` bytes (default is 1024) of the column are used when sorting. You can make more bytes significant in sorting or grouping by increasing its value at server startup or runtime. Clients can change the value of this variable.
- BLOB or TEXT columns in the result of a query that is processed using a temporary table causes the server to use a table on disk rather than in memory because the MEMORY storage engine does not support those data types. Use of disk incurs a performance penalty. Therefore, include BLOB or TEXT columns in the query result only if they are essential.
- BLOB or TEXT types determine the maximum size, but the largest value that can be transmitted between the client and server is determined by the amount of available memory and the size of the communications buffers. Message buffer size can be changed by the `max_allowed_packet` variable, but it must be done for both server and client.

Example

The following example creates a table using a BLOB column with an index.

```
CREATE TABLE test (blob_col BLOB, INDEX(blob_col(10)));
```

For more information, see [CREATE TABLE Statement](#) and [The BLOB and TEXT Types](#) in the *MySQL documentation*.

Oracle synonyms

With Amazon DMS, you can create database objects called synonyms that act as aliases for other schema objects. A synonym is an alternative name for a table, view, sequence, procedure, function, package, materialized view, Java schema object, or other synonym. Synonyms provide data abstraction by hiding the underlying identity of an object, allowing multiple database objects to be referenced by a single name.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Synonyms	Use stored procedures and functions to abstract instance-wide objects.

Oracle usage

Synonyms are database objects that serve as alternative identifiers for other database objects. The referenced database object is called the 'base object' and may reside in the same database, another database on the same instance, or on a remote server.

Synonyms provide an abstraction layer to isolate client application code from changes to the name or location of the base object.

In Oracle, synonyms are often used to simplify the object's name to avoid referring to the other schema as well as for security reasons.

For example, table A resides in schema A, and the client application accesses it through a synonym. Table A needs to be moved to another schema. To make the move seamless, only the synonym definition should be updated. Without synonyms, the client application code must be rewritten to access the other schema or to change the connection string. Instead, you can create a synonym called Table A and it will transparently redirect the calling application to the new schema without any code changes.

You can create synonyms for the following objects:

- Assembly (CLR) stored procedures, table-valued functions, scalar functions, and aggregate functions.
- Stored procedures and functions.
- User-defined tables including local and global temporary tables.
- Views.

Syntax

```
CREATE [OR REPLACE] [EDITIONABLE | NONEDITIONABLE]
[PUBLIC] SYNONYM [schema .] synonym_name
FOR [schema .] object_name [@ dblink];
```

Use the `EDITIONABLE` and `NONEDITIONABLE` options to determine if this object will be private or public. For more information, see [Editioned and Noneditioned Objects](#) in the *Oracle documentation*.

Examples

The following example creates a synonym object `local_emps` that refers to the `usa.emps` table:

```
CREATE SYNONYM local_emps FOR usa.emps;
```

Note

To refer to `local_emps` after you run the preceding command, run your commands or queries against `usa.emps`.

For more information, see [CREATE SYNONYM](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL doesn't support synonyms and there is no known generic workaround.

A partial workaround is to use encapsulating views as an abstraction layer for accessing tables or views. Similarly, you can also use functions or stored procedures that call other functions or stored procedures.

Note

Synonyms are often used in conjunction with Database Links, which are not supported by Aurora MySQL.

For more information, see [MySQL Fully-Qualified Table Names](#), [Views](#), [User-Defined Functions](#), and [Stored Procedures](#).

Oracle and MySQL views

With Amazon DMS, you can create and work with database views in Oracle and MySQL databases. A view is a virtual table that derives its data from one or more underlying tables or views. Views provide a way to present a subset of data from one or more tables, combining data from different tables, or adding additional data transformations.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
			N/A

Oracle usage

Database views store a named SQL query in the Oracle Data Dictionary with a predefined structure. A view doesn't store actual data and may be considered a virtual table or a logical table based on the data from one or more physical database tables.

Privileges

Make sure that the user has the `CREATE VIEW` privilege to create a view in their own schema.

Make sure that the user has the `CREATE ANY VIEW` privilege to create a view in any schema.

Make sure that the owner of the view has all the necessary privileges on the source tables or views on which the view is based (`SELECT` or `DML` privileges).

CREATE (OR REPLACE) VIEW Statements

- `CREATE VIEW` creates a new view.
- `CREATE OR REPLACE` overwrites an existing view and modifies the view definition without having to manually drop and recreate the original view, and without deleting the previously granted privileges.

Oracle common view parameters

Oracle view parameter	Description
<code>CREATE OR REPLACE</code>	Recreate an existing view (if one exists) or create a new view.
<code>FORCE</code>	Create the view regardless of the existence of the source tables or views and regardless of view privileges.
<code>VISIBLE</code> or <code>INVISIBLE</code>	Specify if a column based on the view is visible or invisible.
<code>WITH READ ONLY</code>	Disable DML commands.
<code>WITH CHECK OPTION</code>	Specifies the level of enforcement when performing DML commands on the view.

Examples

Views are classified as either simple or complex.

A *simple view* is a view having a single source table with no aggregate functions. DML operations can be performed on simple views and affect the base table(s). The following example creates and updates a simple View.

```
CREATE OR REPLACE VIEW VW_EMP
AS
SELECT EMPLOYEE_ID, LAST_NAME, EMAIL, SALARY
FROM EMPLOYEES
WHERE DEPARTMENT_ID BETWEEN 100 AND 130;
UPDATE VW_EMP
SET EMAIL=EMAIL||'.org'
WHERE EMPLOYEE_ID=110;
```

```
1 row updated.
```

A *complex view* is a view with several source tables or views containing joins, aggregate (group) functions, or an order by clause. Performing DML operations on complex views can't be done directly, but `INSTEAD OF` triggers can be used as a workaround. The following example creates and updates a complex view.

```
CREATE OR REPLACE VIEW VW_DEP
AS
SELECT B.DEPARTMENT_NAME, COUNT(A.EMPLOYEE_ID) AS CNT
FROM EMPLOYEES A JOIN DEPARTMENTS B USING(DEPARTMENT_ID)
GROUP BY B.DEPARTMENT_NAME;
UPDATE VW_DEP
SET CNT=CNT +1
WHERE DEPARTMENT_NAME=90;

ORA-01732: data manipulation operation not legal on this view
```

For more information, see [CREATE VIEW](#) in the *Oracle documentation*.

MySQL usage

Similar to Oracle, Aurora MySQL views consist of a `SELECT` statement that references base tables and other views.

Aurora MySQL views are created using the `CREATE VIEW` statement. The `SELECT` statement comprising the definition of the view is evaluated only when the view is created and is not affected by subsequent changes to the underlying base tables.

Aurora MySQL views have the following restrictions:

- A view can't reference system variables or user-defined variables.
- When used within a stored procedure or function, the SELECT statement can't reference parameters or local variables.
- A view can't reference prepared statement parameters.
- Make sure that all objects referenced by a view exist when the view is created. If an underlying table or view is later dropped, invoking the view results in an error.
- Views can't reference TEMPORARY tables.
- TEMPORARY views aren't supported.
- Views don't support triggers.
- Aliases are limited to a maximum length of 64 characters and not the typical 256 maximum alias length.

Aurora MySQL provides additional properties that aren't available in Oracle:

- The ALGORITHM clause is a fixed hint that affects the way the MySQL query processor handles the view physical evaluation operator. The MERGE algorithm uses a dynamic approach where the definition of the view is merged to the outer query. The TEMPTABLE algorithm materializes the view data internally. For more information, see [View Processing Algorithms](#) in the *MySQL documentation*.
- You can use the DEFINER and SQL SECURITY clauses can be used to specify a specific security context for checking view permissions at run time.

Similar to Oracle, Aurora MySQL supports updatable views and the ANSI standard CHECK OPTION to limit inserts and updates to rows referenced by the view.

You can use the LOCAL and CASCADED keywords to determine the scope of violation checks. When you use the LOCAL keyword, the CHECK OPTION is evaluated only for the view being created. The CASCADED option causes evaluation of referenced views. The default option is CASCADED.

In general, only views having a one-to-one relationship between the source rows and the exposed rows are updatable. Adding the following constructs prevents modification of data:

- Aggregate functions.
- DISTINCT.

- GROUP BY.
- HAVING.
- UNION or UNION ALL.
- Subquery in the select list.
- Certain joins.
- Reference to a non-updatable view.
- Subquery in the WHERE clause that refers to a table in the FROM clause.
- ALGORITHM = TEMPTABLE.
- Multiple references to any column of a base table.

Make sure that your view has unique column names. Column aliases are derived from the base tables or explicitly specified in the SELECT statement of column definition list. ORDER BY is permitted in Aurora MySQL, but ignored if the outer query has an ORDER BY clause.

A view in Aurora MySQL can invoke functions, which in turn may introduce a change to the database.

Aurora MySQL assesses data access privileges as follows:

- Make sure that the user creating a view has all required privileges to use the top-level objects referenced by the view. For example, for a view referencing table columns, the user must have privilege for each column in the select list of the view definition.
- If the view definition references a stored function, only the privileges needed to invoke the function are checked. The privileges required at run time can be checked only at run time because different invocations may use different execution paths within the function code.
- Make sure that the user referencing a view has the appropriate SELECT, INSERT, UPDATE, or DELETE privileges, as with a normal table.
- When a view is referenced, privileges for all objects accessed by the view are evaluated using the privileges for the view DEFINER account, or the invoker, depending on whether SQL SECURITY is set to DEFINER or INVOKER.
- When a view invocation triggers the execution of a stored function, privileges are checked for statements executed within the function based on the function's SQL SECURITY setting. For functions where the security is set to DEFINER, the function executes with the privileges of the DEFINER account. For functions where it is set to INVOKER, the function executes with the privileges determined by the view's SQL SECURITY setting as described above.

Syntax

```
CREATE [OR REPLACE]
  [ALGORITHM = {UNDEFINED | MERGE | TEMPTABLE}]
  [DEFINER = { <User> | CURRENT_USER }]
  [SQL SECURITY { DEFINER | INVOKER }]
  VIEW <View Name> [(<Column List>)]
  AS <SELECT Statement>
  [WITH [CASCADED | LOCAL] CHECK OPTION];
```

Examples

The following example creates and populate the Invoices table.

```
CREATE TABLE Invoices(
  InvoiceID INT NOT NULL PRIMARY KEY,
  Customer VARCHAR(20) NOT NULL,
  TotalAmount DECIMAL(9,2) NOT NULL);

INSERT INTO Invoices (InvoiceID, Customer, TotalAmount)
VALUES (1, 'John', 1400.23), (2, 'Jeff', 245.00), (3, 'James', 677.22);
```

The following example creates the TotalSales view.

```
CREATE VIEW TotalSales
AS
SELECT Customer, SUM(TotalAmount) AS CustomerTotalAmount
GROUP BY Customer;
```

The following example invokes the view.

```
SELECT * FROM TotalSales
ORDER BY CustomerTotalAmount DESC;

Customer  CustomerTotalAmount
John      1400.23
James     677.22
Jeff      245.00
```

For more information, see [CREATE VIEW Statement](#), [Restrictions on Views](#), and [Updatable and Insertable Views](#) in the *MySQL documentation*.

Oracle XML DB and MySQL XML

With Amazon DMS, you can migrate data between different database engines, including Oracle XML DB and MySQL XML. Oracle XML DB is a feature that provides XML support for storing, processing, and managing XML data in an Oracle database. MySQL XML extends the MySQL server by providing an XML data type for storing XML documents, in addition to functions for extracting and searching XML data.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		XML	Different paradigm and syntax will require application or drivers rewrite.

Oracle usage

Oracle XML DB is a set of Oracle Database technologies providing XML capabilities for database administrators and developers. It provides native XML support and other features including the native XMLType and XMLIndex.

XMLType represents an XML document in the database that is accessible from SQL. It supports standards such as XML Schema, XPath, XQuery, XSLT, and DOM.

XMLIndex supports all forms of XML data from highly structured to completely unstructured.

XML data can be schema-based or non-schema-based. Schema-based XML adheres to an XSD Schema Definition and must be validated. Non-schema-based XML data doesn't require validation.

According to the Oracle documentation, the aspects you should consider when using XML are:

- The ways that you intend to store your XML data.
- The structure of your XML data.
- The languages used to implement your application.
- The ways you intend to process your XML data.

The most common features are:

- **Storage model** — Binary XML.
- **Indexing** — XML search index, XMLIndex with structured component.
- **Database language** — SQL, with SQL/XML functions.
- **XML languages** — XQuery and XSLT.

Storage model — Binary XML

Also called post-parse persistence, it is the default storage model for Oracle XML DB. It is a post-parse, binary format designed specifically for XML data. Binary XML is XML schema-aware and the storage is very flexible.

You can use it for XML schema-based documents or for documents that are not based on an XML schema. You can use it with an XML schema that allows for high data variability or that evolves considerably or unexpectedly.

This storage model also provides efficient partial updating and streaming query evaluation.

The other storage option is object-relational storage and is more efficient when using XML as structured data with a minimum amount of changes and different queries. For more information, see [Oracle XML DB Developer's Guide](#).

Indexing — XML search index, XMLIndex with structured component

XML Search Index provides full-text search over XML data. Oracle recommends storing XMLType data as Binary XML and to use XQuery Full Text (XQFT).

If you are not using binary storage and your data is structured XML, you can use the Oracle text indexes, use the regular string functions such as `contains`, or use XPath `ora:contains`.

If you want to use predicates such as `XMLExists` in your WHERE clause, you must create an XML search index.

Examples

The following example creates a SQL directory object, which is a logical name in the database for a physical directory on the host computer. This directory contains XML files. The example inserts XML content from the `purOrder.xml` file into the `orders` table.

Create an XMLType table.

```
CREATE TABLE orders OF XMLType;
CREATE DIRECTORY xmldir AS path_to_folder_containing_XML_file;
INSERT INTO orders VALUES (XMLType(BFILENAME('XMLDIR',
      'purOrder.xml'),NLS_CHARSET_ID('AL32UTF8')));
```

Create a table with an XMLType column.

```
CREATE TABLE xwarehouses (warehouse_id NUMBER, warehouse_spec XMLTYPE);
```

Create an XMLType view.

```
CREATE VIEW warehouse_view AS
SELECT VALUE(p) AS warehouse_xml FROM xwarehouses p;
```

Insert data into an XMLType column.

```
INSERT INTO xwarehouses
VALUES(100, '<?xml version="1.0"?>
<PO pono="1">
<PNAME>Po_1</PNAME>
<CUSTNAME>John</CUSTNAME>
<SHIPADDR>
<STREET>1033, Main Street</STREET>
<CITY>Sunnyvale</CITY>
<STATE>CA</STATE>
</SHIPADDR></PO>')
```

Create an XML search index and query it with XQuery:

1. After the user gets all the privileges needed and set the right parameter in the Oracle text schema.
2. Create Oracle text section and preference.
3. Create the XML search index (regular index associated with the objects).

```
BEGIN
CTX_DDL.create_section_group('secgroup', 'PATH_SECTION_GROUP');
CTX_DDL.set_sec_grp_attr('secgroup', 'XML_ENABLE', 'T');
CTX_DDL.create_preference('pref', 'BASIC_STORAGE');
CTX_DDL.set_attribute('pref','D_TABLE_CLAUSE', 'TABLESPACE ts_name LOB(DOC) STORE AS
```

```

SECUREFILE(TABLESPACE ts_name COMPRESS MEDIUM CACHE)');
CTX_DDL.set_attribute('pref','I_TABLE_CLAUSE','TABLESPACE ts_name LOB(TOKEN_INFO)
STORE AS SECUREFILE(TABLESPACE ts_name NOCOMPRESS CACHE)');
END;
/
CREATE INDEX po_ctx_idx ON po_binxml(OBJECT_VALUE)
INDEXTYPE IS CTXSYS.CONTEXT
PARAMETERS('storage pref section group secgroup');

```

Query using the preceding index in XQuery. XQuery is W3C standard for generating, querying and updating XML, Natural query language for XML.

Search in the PATH `/PurchaseOrder/LineItems/LineItem/Description` for values containing **Big** and **Street** and then return their **Title** tag (only in the select).

```

SELECT XMLQuery('for $i in /PurchaseOrder/LineItems/LineItem/Description
where $i[.contains text "Big" ftand "Street"] return <Title>{$i}</Title>'
PASSING OBJECT_VALUE RETURNING CONTENT)
FROM po_binxml
WHERE XMlexists('/PurchaseOrder/LineItems/LineItem/Description
[. contains text "Big" ftand "Street"]')

```

XMLIndex with structured component is used for queries that project fixed structured islands of XML content, even if the surrounding data is relatively unstructured. A structured XMLIndex component organizes such islands in a relational format.

Make sure that you define the parts of XML data that you search in queries. This applies to XML schema-based and non-schema-based data.

Create an XMLIndex with a structured component:

1. Create the base XMLIndex on `po_binxml` table. `OBJECT_VALUE` is the XML data stored in the table. All definitions of XML types and Objects are from the XDB schema in the database.
2. Use `DBMS_XMLINDEX.register` parameter to add another structure to the index.
3. Create tables (`po_idx_tab` and `po_index_lineitem`) to store index data as structured data. Next to each table name there is the root of the PATH in the XML data (`/PurchaseOrder` and `/LineItem`). After that, each column is another PATH in this root. Note that in the `po_idx_tab` table the last column is XMLType. It takes everything under this PATH and saves it in XML datatype.
4. Add the group of structure to the index.

```
CREATE INDEX po_xmlindex_ix ON po_binxml (OBJECT_VALUE)
INDEXTYPE IS XDB.XMLIndex PARAMETERS ('PATH TABLE path_tab');
BEGIN
DBMS_XMLINDEX.registerParameter(
'myparam',
'ADD_GROUP GROUP po_item
XMLTable po_idx_tab ''/PurchaseOrder''
COLUMNS reference VARCHAR2(30) PATH ''Reference'',
requestor VARCHAR2(30) PATH ''Requestor'',
username VARCHAR2(30) PATH ''User'',
lineitem XMLType PATH ''LineItems/LineItem'' VIRTUAL
XMLTable po_index_lineitem ''/LineItem'' PASSING lineitem
COLUMNS itemno BINARY_DOUBLE PATH ''@ItemNumber'',
description VARCHAR2(256) PATH ''Description'',
partno VARCHAR2(14) PATH ''Part/@Id'',
quantity BINARY_DOUBLE PATH ''Part/@Quantity'',
unitprice BINARY_DOUBLE PATH ''Part/@UnitPrice''');
END;
/

ALTER INDEX po_xmlindex_ix PARAMETERS('PARAM myparam');
```

For more information, see [Indexes for XMLType Data](#) in the *Oracle documentation*.

SQL/XML functions

Oracle Database provides two main SQL/XML groups:

- SQL/XML publishing functions.
- SQL/XML query and update functions.

SQL/XML publishing functions

SQL/XML publishing functions are SQL results generated from XML data. They are also called SQL/XML generation functions.

XMLQuery is used in SELECT clauses to return the result as XMLType data. See the previous example for creating an XML search index.

XMLTable is used in FROM clauses to get results using XQuery, and insert the results into a virtual table. This function can insert data into existing database table.

```

SELECT po.reference, li.*
FROM po_binaryxml p,
XMLTable('/PurchaseOrder' PASSING p.OBJECT_VALUE
COLUMNS
reference VARCHAR2(30) PATH 'Reference',
lineitem XMLType PATH 'LineItems/LineItem') po,
XMLTable('/LineItem' PASSING po.lineitem
COLUMNS
itemno NUMBER(38) PATH '@ItemNumber',
description VARCHAR2(256) PATH 'Description',
partno VARCHAR2(14) PATH 'Part/@Id',
quantity NUMBER(12, 2) PATH 'Part/@Quantity',
unitprice NUMBER(8, 4) PATH 'Part/@UnitPrice') li;

```

XMLExists is used in WHERE clauses to check if an XQuery expression returns a non-empty query sequence. If it does, it returns TRUE. Otherwise, it returns FALSE. In the following example, the query searches the purchaseorder table for PurchaseOrders that where the SpecialInstructions tag is set to Expedite.

```

SELECT OBJECT_VALUE FROM purchaseorder
WHERE XMLExists('/PurchaseOrder[SpecialInstructions="Expedite"]'
PASSING OBJECT_VALUE);

```

XMLCast is used in SELECT clauses to convert scalar values returned from XQuery to NUMBER, VARCHAR2, CHAR, CLOB, BLOB, REF, or XMLType. For example, after finding the objects that have SpecialInstructions set to Expedite, XMLCast returns the Reference in each item as VARCHAR2(100).

```

SELECT XMLCast(XMLQuery('/PurchaseOrder/Reference'
PASSING OBJECT_VALUE
RETURNING CONTENT) AS VARCHAR2(100)) "REFERENCE"
FROM purchaseorder
WHERE XMLExists('/PurchaseOrder[SpecialInstructions="Expedite"]'
PASSING OBJECT_VALUE);

```

For more information, see [XMLELEMENT](#) in the *Oracle documentation*.

SQL/XML query and update functions

SQL/XML query and update functions are used to query and update XML content as part of regular SQL operations.

For XMLQuery, see the example preceding.

In the following example, after finding the relevant item with `XMLExists` in the set clause, the command sets the `OBJECT_VALUE` to a new `NEW-DAUSTIN-20021009123335811PDT.xml` file located in the `XMLDIR` directory.

```
UPDATE purchaseorder po
SET po.OBJECT_VALUE = XMLType(bfilename('XMLDIR', 'NEW-
DAUSTIN-20021009123335811PDT.xml'),
    nls_charset_id('AL32UTF8'))
WHERE XMLExists('$p/PurchaseOrder[Reference="DAUSTIN-20021009123335811PDT"]'
    PASSING po.OBJECT_VALUE AS "p");
```

For more information, see [XMLQUERY](#) in the *Oracle documentation*.

SQL and PL/SQL

Conversion of SQL and PL/SQL is covered in the [SQL and PL/SQL](#) topic.

MySQL usage

Aurora MySQL support for unstructured data is the opposite of Oracle. There is minimal support for XML, but a native JSON data type and more than 25 dedicated JSON functions.

XML support

Aurora MySQL supports two XML functions: `ExtractValue` and `UpdateXML`.

`ExtractValue` accepts an XML document, or fragment, and an XPATH expression. The function returns the character data of the child or element matched by the XPATH expression. If there is more than one match, the function returns the content of child nodes as a space delimited character string. `ExtractValue` returns only CDATA and doesn't return tags and sub-tags contained within a matching tag or its content.

Consider the following example.

```
SELECT ExtractValue('<Root><Person>John</Person>
<Person>Jim</Person></Root>', '/Root/Person');
```

For the preceding example, the result looks as shown following.

```
John Jim
```

You can use `UpdateXML` to replace an XML fragment with another fragment using XPATH expressions similar to `ExtractValue`. If a match is found, it returns the new, updated XML. If there are no matches, or multiple matches, the original XML is returned.

Consider the following example.

```
SELECT UpdateXML('<Root><Person>John</Person>
<Person>Jim</Person></Root>', '/Root', '<Person>Jack</Person>')
```

For the preceding example, the result looks as shown following.

```
<Person>Jack</Person>
```

Note

Aurora MySQL doesn't support MySQL `LOAD XML` syntax. For more information, see [Loading data into an Aurora MySQL DB cluster from text files in an Amazon S3 bucket](#) in the *User Guide for Aurora*.

Summary

Description	Oracle	Aurora MySQL
XML functions	XMLQuery, XPath, XMLTable, XMLEExists, and XMLCast	ExtractValue and UpdateXML
Create a table with XML	CREATE TABLE test OF XMLType; or CREATE TABLE test (doc XMLType);	Not supported
Insert data into xml column	<pre>INSERT INTO test VALUES ('<?xml version="1.0"?></pre>	XML data can be loaded into regular tables from S3. For more information, see

Description	Oracle	Aurora MySQL
	<pre><PO pono="1"> <PNAME>Po _1</PNAME> <CUSTNAME>John</ CUSTNAME> <SHIPADDR> <STREET>1033, Main Street</STREET> <CITY>Sunnyvale</C ITY> <STATE>CA</STATE> </SHIPADDR> </PO>')</pre>	<p>Loading data into an Aurora MySQL DB cluster from text files in an Amazon S3 bucket in the <i>User Guide for Aurora</i>.</p>
Create Index	<pre>CREATE INDEX test_idx ON test (OBJECT_VALUE) INDEXTYPE IS XDB.XMLIn dex PARAMETERS ('PATH TABLE path_tab'); BEGIN DBMS_XMLINDEX.reg isterParameter('myparam', 'ADD_GROUP GROUP a_item XMLTable test_idx_tab ''/Path'' COLUMNS tag VARCHAR2(30) PATH 'tag''); END; / ALTER INDEX test_idx PARAMETERS ('PARAM myparam');</pre>	<p>Requires adding always generated computed and persisted columns with JSON expressions and indexing them explicitly. The optimizer can make use of JSON expressions only.</p>

Description	Oracle	Aurora MySQL
Create a full-text index	<p>After preference and section created in Oracle Text</p> <pre>CREATE INDEX test_idx ON test (OBJECT_VALUE) INDEXTYPE IS CTXSYS.CO NTEXT PARAMETERS('storage pref section group secgroup');</pre>	N/A
Query using XQuery	<pre>SELECT XMLQuery('for \$i in /PurchaseOrder/ LineItems/LineItem/ Description where \$i[. contains text "Big"] return <Title>{\$i}</ Title>' PASSING OBJECT_VALUE RETURNING CONTENT) FROM xml_tbl;</pre>	N/A
Query using XPath	<pre>select sys.XMLType pe.extract (doc,'/student/f irstname/text()') firstname from test;</pre>	Because there is no XML data type, doc uses VARCHAR to store the XML content [source] ---- select ExtractValue (doc,'//student//firstn ame') firstname from test; ----

Description	Oracle	Aurora MySQL
Function to check if tag exists and function to cast and return a string data type	<pre>SELECT XMLCast(XMLQuery ('/PurchaseOrder/Reference' PASSING OBJECT_VALUE RETURNING CONTENT) AS VARCHAR2(100)) "REFERENCE" FROM purchaseorder WHERE XMLElementExists ('/PurchaseOrder[SpecialInstructions ="Expedite"]' PASSING OBJECT_VALUE);</pre>	N/A
Validate schema using XSD	Supported	Not supported

For more information, see [XML Functions](#) in the *MySQL documentation*.

Oracle table compression

With Amazon DMS, you can optimize storage utilization and improve query performance for Oracle databases by leveraging table compression. Oracle table compression reduces the disk space footprint of tables and associated indexes, which can lead to significant cost savings, especially for large datasets.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Syntax and option differences, similar functionality. MySQL doesn't compress partitions.

Oracle usage

Oracle table compression reduces the size of data. It saves disk space, reduces memory usage, and speeds up query execution during reads. However, the cost is increased CPU overhead for data loading and DML.

Table compression is completely transparent to applications. It is most commonly used for OLAP systems where there are significantly more read operations, but it can also be used in OLTP systems.

Tables can be compressed when they are created using the `COMPRESS` clause. Existing tables can be compressed using the `COMPRESS` clause with an `ALTER TABLE` statement.

You can turn on compression for `ALL OPERATIONS` on the table or for `DIRECT_LOAD OPERATIONS` only. When compression is turned on for all operations, compression occurs during all DML statements and when data is inserted with a bulk (direct-path) insert operation.

The compression clause provides four options:

- `NOCOMPRESS` — Don't use compression. This is the default option.
- `COMPRESS` — Turns on compression on the table or partition during direct-path inserts only.
- `COMPRESS FOR DIRECT_LOAD OPERATIONS` — Turns on compression on the table or partition during direct-path inserts only.
- `COMPRESS FOR ALL OPERATIONS` — Turns on the compression for all operations including DML statements. This option is mostly used for OLTP systems.

Examples

View the compression status of tables.

```
SELECT OWNER, TABLE_NAME, COMPRESSION, COMPRESS_FOR FROM dba_tables;
```

The following example creates a compressed table.

```
CREATE TABLE comp_tbl  
(id NUMBER NOT NULL,  
created_date DATE NOT NULL)  
COMPRESS FOR ALL OPERATIONS;
```

The following example creates a partitioned table with a compressed partition.

```
CREATE TABLE comp_part_tbl
(id NUMBER NOT NULL,
created_date DATE NOT NULL)
PARTITION BY RANGE (created_date) (
PARTITION comp_part_tbl_q1 VALUES LESS THAN (TO_DATE('01/01/2018', 'DD/MM/YYYY'))
COMPRESS,
PARTITION comp_part_tbl_q2 VALUES LESS THAN (TO_DATE('01/04/2018', 'DD/MM/YYYY'))
COMPRESS FOR DIRECT_LOAD OPERATIONS,
PARTITION comp_part_tbl_q3 VALUES LESS THAN (TO_DATE('01/07/2018', 'DD/MM/YYYY'))
COMPRESS FOR ALL OPERATIONS,
PARTITION comp_part_tbl_q4 VALUES LESS THAN (MAXVALUE) NOCOMPRESS);
```

For more information, see [DBMS_COMPRESSION](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL doesn't support compressed tables (that is, tables created with `ROW_FORMAT=COMPRESSED`). Make sure that you expand your compressed tables by setting `ROW_FORMAT` to `DEFAULT`, `COMPACT`, `DYNAMIC`, or `REDUNDANT`.

For more information, see [InnoDB Table Compression](#) in the *MySQL documentation*.

Oracle Log Miner and MySQL logs

With Amazon DMS, you can capture data manipulation language (DML) operations for replication or auditing purposes using Oracle Log Miner and MySQL binary logs. Oracle Log Miner provides access to redo log files, enabling the reconstruction and analysis of database activity. MySQL binary logs record all statements that update data or potentially could have updated it.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	MySQL doesn't support LogMiner, workaround is available.

Oracle usage

Oracle Log Miner is a tool for querying the database Redo Logs and the Archived Redo Logs using an SQL interface. Using Log Miner, you can analyze the content of database transaction logs (online and archived redo logs) and gain historical insights on past database activity such as data modification by individual DML statements.

Examples

The following examples demonstrate how to use Log Miner to view DML statements that run on the employees table.

Find the current redo log file.

```
SELECT V$LOG.STATUS, MEMBER
FROM V$LOG, V$LOGFILE
WHERE V$LOG.GROUP# = V$LOGFILE.GROUP#
AND V$LOG.STATUS = 'CURRENT';
```

STATUS	MEMBER
CURRENT	/u01/app/oracle/oradata/orcl/redo02.log

Use the `DBMS_LOGMNR.ADD_LOGFILE` procedure. Pass the file path as a parameter to the Log Miner API.

```
BEGIN
DBMS_LOGMNR.ADD_LOGFILE('/u01/app/oracle/oradata/orcl/redo02.log');
END;
/
```

PL/SQL procedure successfully completed.

Start Log Miner using the `DBMS_LOGMNR.START_LOGMNR` procedure.

```
BEGIN
DBMS_LOGMNR.START_LOGMNR(options=>
dbms_logmnr.dict_from_online_catalog);
END;
/
```

PL/SQL procedure successfully completed.

Run a DML statement.

```
UPDATE HR.EMPLOYEES SET SALARY=SALARY+1000 WHERE EMPLOYEE_ID=116;
COMMIT;
```

Query the V\$LOGMNR_CONTENTS table to view the DML commands captured by the Log Miner.

```
SELECT TO_CHAR(TIMESTAMP, 'mm/dd/yy hh24:mi:ss') TIMESTAMP,
SEG_NAME, OPERATION, SQL_REDO, SQL_UNDO
FROM V$LOGMNR_CONTENTS
WHERE TABLE_NAME = 'EMPLOYEES'
AND OPERATION = 'UPDATE';

TIMESTAMP    SEG_NAME    OPERATION
10/09/17     06:43:44   EMPLOYEES  UPDATE

SQL_REDO                                           SQL_UNDO
update "HR"."EMPLOYEES" set
"SALARY" = '3900' where "SALARY" = '2900'
'3900'
and ROWID = 'AAAViUAAEAAABVvAAQ';                and ROWID = 'AAAViUAAEAAABVvAAQ';
```

For more information, see [Using LogMiner to Analyze Redo Log Files](#) in the *Oracle documentation*.

MySQL usage

The mysqlbinlog utility is the MySQL equivalent to Oracle Log Miner. You can use Log Miner to search for many types of information. This topic covers all of the MySQL logs that are available so you can decide which log is best for your use case.

Aurora MySQL generates four logs that can be viewed by database administrators:

- **Error log** — Contains information about errors and server start and stop events.
- **General query log** — Contains a general record of MySQL operations such as connect, disconnect, queries, and so on.
- **Slow query log** — Contains a log of slow SQL statements.
- **Bin log** — When used, contains row and statement levels of commands records.

The MySQL error log is generated by default. You can generate the slow query and general logs by setting parameters in the database parameter group. Amazon RDS rotates all MySQL log files.

You can monitor the MySQL logs directly through the Amazon RDS console, Amazon RDS API, Amazon CLI, or Amazon SDKs. You can also access MySQL logs by directing the logs to a database table in the main database and then querying that table. You can use the `mysqlbinlog` utility to download a binary log.

Downloading MySQL binlog files

The binlog in MySQL is used for replication needs. MySQL uses it to replicate commands between master MySQL server to slave server. These logs can be read using the `mysqlbinlog` utility.

The `mysqlbinlog` utility is equivalent to Oracle Log Miner and enables users to read the server's binary log (similar to the Oracle redo log). The server's binary log consists of files that describe modifications to database contents (events).

While these logs do not contain a lot of information, they can provide needed data for some use cases.

To download and read the binary log, check to see if the binlog is activated by typing this command:

```
SHOW BINARY LOGS;
```

Note

If the binlog isn't activated, this command returns an error. If the binlog is activated, the binlog files list is displayed.

After querying the binlog files list you can select a file to download by using this command:

```
mysqlbinlog
  --read-from-remote-server
  --host=mysql-cluster1.cluster-crqdlsqlnpry.useast-1.rds.amazonaws.com
  --port=3306
  --user naya
  --password mysql-bin-changelog.0098
```

For more information, see [MySQL database log files](#) in the *Amazon Relational Database Service User Guide*.

The output example for binlog looks as shown following:

```

use `aws`/*!*/;
SET TIMESTAMP=1551125550/*!*/;
SET @@session.pseudo_thread_id=12/*!*/;
SET @@session.foreign_key_checks=1, @@session.sql_auto_is_null=0, @@session.unique_
checks=1, @@session.autocommit=1/*!*/;
SET @@session.sql_mode=2097152/*!*/;
SET @@session.auto_increment_increment=1, @@session.auto_increment_offset=1/*!*/;
/*!\C utf8 *//*!*/;
SET @@session.character_set_client=33,@@session.collation_connection=
33,@@session.collation_server=8/*!*/;
SET @@session.lc_time_names=0/*!*/;
SET @@session.collation_database=DEFAULT/*!*/;
last_committed=1 sequence_number=2 rbr_only=no original_committed_
timestamp=0 immediate_commit_timestamp=0 transaction_length=0
# original_commit_timestamp=0 (1969-12-31 19:00:00.000000 Eastern Standard Time)
# immediate_commit_timestamp=0 (1969-12-31 19:00:00.000000 Eastern Standard Time)
/*!80001 SET @@session.original_commit_timestamp=0/*!*/;
/*!80014 SET @@session.original_server_version=0/*!*/;
/*!80014 SET @@session.immediate_server_version=0/*!*/;
SET @@SESSION.GTID_NEXT= 'ANONYMOUS'/*!*/;
# at 434
#190225 15:12:50 server id 565151648 end_log_pos 513 CRC32 0x1188c639 Query
thread_id=12 exec_time=0 error_code=0
SET TIMESTAMP=1551125570/*!*/;
BEGIN
/*!*/;
# at 513
#190225 15:12:50 server id 565151648 end_log_pos 669 CRC32 0x051c3800 Query
thread_id=12 exec_time=0 error_code=0
SET TIMESTAMP=1551125570/*!*/;
/* ApplicationName=mysql */ insert into test values (1),(1),(1)
/*!*/;
# at 669
#190225 15:12:50 server id 565151648 end_log_pos 700 CRC32 0x72697ff4 Xid = 5467
COMMIT/*!*/;
SET @@SESSION.GTID_NEXT= 'AUTOMATIC' /* added by mysqlbinlog */ /*!*/;
DELIMITER ;
# End of log file

```

For more information, see [mysqlbinlog — Utility for Processing Binary Log Files](#) in the *MySQL documentation*.

Accessing MySQL error logs

The MySQL error log is written to the `mysql-error.log` file. You can view `mysql-error.log` by using the Amazon RDS console or by retrieving the log using the Amazon RDS API, Amazon RDS CLI, or Amazon SDKs. `mysqlerror.log` is flushed every 5 minutes and its contents are appended to `mysql-error-running.log`. The `mysql-errorrunning.log` file is then rotated every hour. The hourly files generated during the last 24 hours are retained. Each log file has the hour it was generated (in UTC) appended to its name. The log files also have a timestamp that helps you determine when the log entries were written.

MySQL writes to the error log only on startup, shutdown, and when it encounters errors. A database instance can go hours or days without new entries being written to the error log. If you see no recent entries, it's because the server did not encounter an error that would result in a log entry.

Accessing the MySQL slow query and general logs

The MySQL slow query log and the general log can be written to a file or a database table by setting parameters in the database parameter group. You must set these parameters before you can view the slow query log or general log in the Amazon RDS console, Amazon RDS API, Amazon RDS CLI, or Amazon SDKs.

You can control MySQL logging by using the following parameters:

- `slow_query_log` — To create the slow query log, set to 1. The default is 0.
- `general_log` — To create the general log, set to 1. The default is 0.
- `long_query_time` — To prevent fast-running queries from being logged in the slow query log, specify a value for the shortest query run time in seconds to be logged. The default is 10 seconds; the minimum is 0. If `log_output = FILE`, you can specify a floating point value with a resolution of microseconds. If `log_output = TABLE`, make sure that you specify an integer value with a resolution of seconds. Only queries where the execution time exceeds the `long_query_time` value are logged. For example, setting `long_query_time` to 0.1 prevents a query that runs for less than 100 milliseconds from being logged.
- `log_queries_not_using_indexes` — To log all queries that do not use an index to the slow query log, set to 1. The default is 0. Queries that do not use an index are logged even if their execution time is less than the value of the `long_query_time` parameter.
- `log_output` — You can specify one of the following options for the `log_output` parameter.

- **TABLE** — Write general queries to the `mysql.general_log` table, and write slow queries to the `mysql.slow_log` table. This is the default option.
- **FILE** — Write both general and slow query logs to the file system. Log files are rotated hourly.
- **NONE** — Turn off logging.

You can configure a MySQL instance to publish log data to a log group in Amazon CloudWatch Logs. CloudWatch Logs support real-time analysis of the log data, create alarms, and view metrics. You can use CloudWatch Logs to store your log records in highly durable storage. For more information, see [MySQL Database Log Files](#) in the *Amazon Relational Database Service User Guide*.

Amazon RDS normally purges a binary log as soon as possible, but the binary log must still be available on the instance to be accessed by `mysqlbinlog`. To specify the number of hours for RDS to retain binary logs, use the `mysql.rds_set_configuration` stored procedure and specify a period with enough time for you to download the logs. After you set the retention period, monitor storage usage for the database instance to ensure the retained binary logs don't consume too much storage.

Examples

Determine the output location of the logs and if slow query and general logging are turned on.

```
select @@GLOBAL.log_output, @@GLOBAL.slow_query_log, @@GLOBAL.general_log
```

To view the logs using Amazon Management Console:

1. Sign in to the Amazon Management Console and choose **RDS**.
2. Choose your DB instance and scroll down to the **Logs** section.
3. Choose a log to inspect or download.

The following example configures retention of the binary logs (in hours). In this example the binary log will be retained one day.

```
call mysql.rds_set_configuration('binlog retention hours', 24);
```

For more information, see [The Binary Log](#) in the *MySQL documentation* and [MySQL database log files](#) in the *Amazon Relational Database Service User Guide*.

Oracle SQL Result Cache and MySQL Query Cache

With Amazon DMS, you can leverage performance optimization features such as Oracle SQL Result Cache and MySQL Query Cache to improve query execution times. The Oracle SQL Result Cache stores data from previous queries, allowing faster retrieval for identical subsequent queries. MySQL Query Cache temporarily stores the text of a SELECT query and its corresponding result set, facilitating quicker responses to repeated queries on the same data.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	Syntax and option differences, similar functionality. This is off the MySQL roadmap and suggested not to be used.

Oracle usage

The Oracle SQL Result Cache feature is related to the following caching categories:

- Global temporary tables.
- Materialized views.
- PL/SQL collection.
- The WHEN clause.

The Result Cache reduces I/O operations by skipping the fetch step of execution plans and retrieving rows from the buffer cache. This feature is most useful for data warehouse scenarios where many rows must be scanned, but the result sets contain few rows. The rows are stored in the System Global Area (SGA) and are reused when the same SQL statements are executed in the current session or other sessions.

The `RESULT_CACHE_MODE` parameter controls caching and accepts the following values:

- **MANUAL** — SQL results are not cached for SQL statements unless they use a hint to perform caching.
- **FORCE** — All results are cached for SQL statements unless they use a hint to prevent caching.

In Oracle Real Application Cluster (RAC) environments, each instance has its own private result cache and can't be used by other instances.

The query result cache is not compatible with scalar subquery caching.

Examples

Cache a query when `RESULT_CACHE_MODE` is set to `MANUAL`.

```
SELECT /*+ RESULT_CACHE */ count(*) FROM bigdata_smallres_tbl;
```

Turn off caching when `RESULT_CACHE_MODE` is set to `FORCE` and a result cache isn't needed.

```
SELECT /*+ NO_RESULT_CACHE */ count(*) FROM bigdata_smallres_tbl;
```

For more information, see [Configuring the Client Result Cache](#) in the *Oracle documentation*.

MySQL usage

According to the MySQL roadmap, it is recommended not to use the Query Cache.

Like the Oracle Result Cache, the MySQL Query Cache reduces I/O operations by skipping the fetch step of run plans and retrieving rows from the buffer cache. It can be shared across multiple sessions.

The Query Cache is deprecated as of MySQL 5.7.20 and will be removed in MySQL 8.0. For more information, see [Retiring Support for the Query Cache](#) in the *MySQL Blog*.

Examples

The following example runs a `select` statement using the Query Cache.

```
SELECT SQL_CACHE count(*) FROM bigdata_smallres_tbl;
```

The following example runs a `select` statement without using the Query Cache.

```
SELECT SQL_NO_CACHE count(*) FROM bigdata_smallres_tbl;
```

For more information, see [The MySQL Query Cache](#) in the *MySQL documentation*.

Oracle and MySQL high availability and disaster recovery

This section includes pages about Oracle and MySQL high availability and disaster recovery capabilities.

Topics

- [Oracle Active Data Guard and MySQL replicas](#)
- [Oracle Real Application Clusters and Aurora MySQL architecture](#)
- [Migrate to Aurora MySQL Serverless](#)
- [Oracle Traffic Director and Amazon RDS Proxy for Amazon Aurora MySQL](#)
- [Oracle Data Pump and MySQL mysqldump and mysql](#)
- [Oracle Flashback Database and MySQL snapshots](#)
- [Oracle Flashback Table and MySQL snapshots](#)
- [Oracle Recovery Manager and Amazon RDS snapshots](#)
- [Oracle SQL*Loader and MySQL mysqlimport and LOAD DATA](#)

Oracle Active Data Guard and MySQL replicas


With Amazon DMS, you can create and manage Oracle Active Data Guard and MySQL replicas to achieve high availability and data redundancy for your databases. Oracle Active Data Guard provides a physical standby database that remains synchronized with the primary database, while MySQL replicas maintain an identical copy of data from a source database instance.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Distribute load, applications, or users across multiple instances.

Oracle usage

Oracle Active Data Guard (ADG) is a synced database architecture with primary and standby databases. The difference between Data Guard and ADG is that ADG standby databases allow read access only.

The following diagram illustrates the ADG architecture.



- **Primary DB** — The main database open to read and write operations.
- **Redo/Archive** — The redo files and archives that store the redo entries for recovery operations.
- **Data Broker** — The data guard broker service is responsible for all failover and syncing operations.
- **Standby DB** — The secondary database that allows read operations only. This database remains in recovery mode until it is shut down or becomes the primary (failover or switchover).
- **Log Apply** — Runs all the redo log entries from the redo and archives files on the standby db.

- **Redo/Archive** — Contains the redo files and archives that are synced from the primary log and archive files.
- **Data Broker** — The Data Guard broker service is responsible for all failover and syncing operations.

All components use SQL*NET protocol.

Special features

- You can select asynchronously for best performance or synchronously for best data protection.
- You can temporarily convert a standby database to a snapshot database and allow read/write operations. When you are done running QA, testing, loads, or other operations, it can be switched back to standby.
- A sync gap can be specified between the primary and standby databases to account for human errors (for example, creating 12 hours gap of sync).

For more information, see [Creating a Physical Standby Database](#) in the *Oracle documentation*.

MySQL usage

You can use Aurora replicas for scaling read operations and increasing availability such as Oracle Active Data Guard, but with less configuration and administration. You can easily manage many replicas from the Amazon RDS console. Alternatively, you can use the Amazon CLI for automation.

When you create Aurora MySQL instances, use one of the two following replication options:

- **Multi-AZ (Availability Zone)** — Create a replicating instance in a different region.
- **Instance Read Replicas** — Create a replicating instance in the same region.

For instance options, you can use one of the two following options:

- Create Aurora Replica.
- Create Cross Region Read Replica.

The main differences between these two options are:

- Cross Region creates a new reader cluster in a different region. Use Cross Region for a higher level of Higher Availability and to keep data closer to the end users.
- Cross Region has more lag between the two instances.
- Additional charges apply for transferring data between two regions.

To view the most current lag value between the primary and replicas, query the `mysql.ro_replica_status` table and check the `Replica_lag_in_msec` column. This column value is provided to Amazon CloudWatch as the `ReplicaLag` metric. The values in the `mysql.ro_replica_status` are also provided in the `INFORMATION_SCHEMA.REPLICA_HOST_STATUS` table in your Aurora MySQL DB cluster.

DDL statements that run on the primary instance may interrupt database connections on the associated Aurora Replicas. If an Aurora Replica connection is actively using a database object such as a table, and that object is modified on the primary instance using a DDL statement, the Aurora Replica connection is interrupted.

Rebooting the primary instance of an Amazon Aurora database cluster also automatically reboots the Aurora Replicas for that database cluster.

Before you create a cross region replica, turn on the `binlog_format` parameter.

When using Multi-AZ, the primary database instance switches over automatically to the standby replica if any of the following conditions occur:

- The primary database instance fails.
- An Availability Zone outage.
- The database instance server type is changed.
- The operating system of the database instance is undergoing software patching.
- A manual failover of the database instance was initiated using `reboot with failover`.

Examples

The following walkthrough demonstrates how to create a replica reader.

1. Sign in to your Amazon console and choose **RDS**.
2. Choose **Instance actions** and choose **Create cross-Region read replica**.
3. Enter all required details and choose **Create**.

After the replica is created, you can run read and write operations on the primary instance and read-only operations on the replica.

For more information, see [Single-master replication with Amazon Aurora MySQL](#), [Replicating Amazon Aurora MySQL DB clusters](#), and [Creating an Amazon Aurora DB cluster](#) in the *User Guide for Aurora*.

Oracle Real Application Clusters and Aurora MySQL architecture


With Amazon DMS, you can migrate Oracle Real Application Clusters (RAC) and Aurora MySQL databases to the Amazon Cloud. Oracle RAC is a database clustering solution that provides high availability, scalability, and load balancing for Oracle databases. Aurora MySQL is a fully managed MySQL-compatible relational database service offered by Amazon. Migrating these databases can be beneficial for organizations seeking to reduce operational overhead, improve performance, and leverage the scalability and reliability of Amazon services.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Distribute load, applications, or users across multiple instances.

Oracle usage

Oracle Real Application Clusters (RAC) is one of the most advanced and capable technologies providing highly available and scalable relational databases. It allows multiple Oracle instances to access a single database. Applications can access the database through the multiple instances in Active-Active mode.

The following diagram illustrates the Oracle RAC architecture.



Oracle RAC requires network configuration of SCAN IPs, VIP IPs, interconnect, and other items. As a best practice, all servers should run the same versions of Oracle software.

Because of the shared nature of the RAC cluster architecture—specifically, having all nodes write to a single set of database data files on disk—the following two special coordination mechanisms ensure Oracle database objects and data maintain ACID compliance:

- **Global Cache Services (GCS)** — Tracks the location and status of the database data blocks and helps guarantee data integrity for global access across all cluster nodes.
- **Global Enqueue Services (GES)** — Performs concurrency control across all cluster nodes including cache locks and transactions.

These services, which run as background processes on each cluster node, are essential for serializing access to shared data structures in an Oracle database.

Shared storage is another essential component in the Oracle RAC architecture. All cluster nodes read and write data to the same physical database files stored on a disk accessible by all nodes. Most customers rely on high-end storage hardware to provide the shared storage capabilities required for RAC.

In addition, Oracle provides its own software-based storage/disk management mechanism called Automatic Storage Management (ASM). ASM is implemented as a set of special background processes that run on all cluster nodes and allow for easy management of the database storage layer.

Performance and Scale-Out with Oracle RAC

You can add new nodes to an existing RAC cluster without downtime. Adding more nodes increases the level of high availability and enhances performance.

Although you can scale read performance easily by adding more cluster nodes, scaling write performance is more complicated. Technically, Oracle RAC can scale writes and reads together when adding new nodes to the cluster, but attempts from multiple sessions to modify rows that reside in the same physical Oracle block (the lowest level of logical I/O performed by the database) can cause write overhead for the requested block and impact write performance.

Concurrency is another reason why RAC implements a “smart mastering” mechanism that attempts to reduce write-concurrency overhead. The “smart mastering” mechanism enables the database to determine which service causes which rows to be read into the buffer cache and master the data blocks only on those nodes where the service is active. Scaling writes in RAC isn’t as straightforward as scaling reads.

With the limitations for pure write scale-out, many Oracle RAC customers choose to split their RAC clusters into multiple services, which are logical groupings of nodes in the same RAC cluster. By using services, you can use Oracle RAC to perform direct writes to specific cluster nodes. This is usually done in one of two ways:

- Splitting writes from different individual modules in the application (that is, groups of independent tables) to different nodes in the cluster. This approach is also known as application partitioning (not to be confused with database table partitions).
- In extremely non-optimized workloads with high concurrency, directing all writes to a single RAC node and load-balancing only the reads.

In summary, Oracle Real Application Clusters provides two major benefits:

- Multiple database nodes within a single RAC cluster provide increased high availability. No single point of failure exists from the database servers themselves. However, the shared storage requires storage-based high availability or disaster recovery solutions.
- Multiple cluster database nodes enable scaling-out query performance across multiple servers.

For more information, see [Oracle Real Application Clusters](#) in the *Oracle documentation*.

MySQL usage

Aurora extends the vanilla versions of MySQL in two major ways:

- Adds enhancements to the MySQL database kernel itself to improve performance (concurrency, locking, multi-threading, and so on).
- Uses the capabilities of the Amazon ecosystem for greater high availability, disaster recovery, and backup or recovery functionality.

Comparing the Amazon Aurora architecture to Oracle RAC, there are major differences in how Amazon implements scalability and increased high availability. These differences are due mainly to the existing capabilities of MySQL and the strengths the Amazon backend provides in terms of networking and storage.


Instead of having multiple read/write cluster nodes access a shared disk, an Aurora cluster has a single primary node that is open for reads and writes and a set of replica nodes that are open for reads with automatic promotion to primary in case of failures. While Oracle RAC uses a set of background processes to coordinate writes across all cluster nodes, the Amazon Aurora primary writes a constant redo stream to six storage nodes distributed across three Availability Zones within an Amazon Region. The only writes that cross the network are redo log records, not pages.

Each Aurora cluster can have one or more instances serving different purposes:

- At any given time, a single instance functions as the primary that handles both writes and reads from your applications.
- You can create up to 15 read replicas in addition to the primary, which are used for two purposes:
 - **Performance and Read Scalability** — Replicas can be used as read-only nodes for queries and report workloads.

- **High Availability** — Replicas can be used as failover nodes in the event the master fails. Each read replica can be located in one of the three Availability Zones hosting the Aurora cluster. A single Availability Zone can host more than one read replica.

The following diagram illustrates a high-level Aurora architecture with four cluster nodes: one primary and three read replicas. The primary node is located in Availability Zone A, the first read replica in Availability Zone B, and the second and third read replicas in Availability Zone C.




An Aurora Storage volume is made up of 10 GB segments of data with six copies spread across three Availability Zones. Each Amazon Aurora read replica shares the same underlying volume as the master instance. Updates made by the master are visible to all read replicas through a combination of reading from the shared Aurora storage volume and applying log updates in-memory when received from the primary instance after a master failure. Promotion of a read replica to master usually occurs in less than 30 seconds with no data loss.

For a write to be considered durable in Aurora, the primary instance sends a redo stream to six storage nodes — two in each availability zone for the storage volume — and waits until four of the six nodes have responded. No database pages are ever written from the database tier to the storage tier. The Aurora Storage volume asynchronously applies redo records to generate database pages in the background or on demand. Aurora hides the underlying complexity.

High availability and scale-out in Aurora

Aurora provides two endpoints for cluster access. These endpoints provide both high availability capabilities and scale-out read processing for connecting applications.

- **Cluster endpoint** — Connects to the current primary instance for the Aurora cluster. You can perform both read and write operations using the cluster endpoint. If the current primary instance fails, Aurora automatically fails over to a new primary instance. During a failover, the database cluster continues to serve connection requests to the cluster endpoint from the new primary instance with minimal interruption of service.
- **Reader endpoint** — Provides load-balancing capabilities (round-robin) across the replicas allowing applications to scale-out reads across the Aurora cluster. Using the Reader Endpoint provides better use of the resources available in the cluster. The reader endpoint also enhances high availability. If an Amazon Availability Zone fails, the application's use of the reader endpoint continues to send read traffic to the other replicas with minimal disruption.



While Amazon Aurora focuses on the scale-out of reads and Oracle RAC can scale-out both reads and writes, most OLTP applications are usually not limited by write scalability. Many Oracle RAC customers use RAC first for high availability and second to scale-out their reads. You can write to any node in an Oracle RAC cluster, but this capability is often a functional benefit for the application versus a method for achieving unlimited scalability for writes.


Summary

- In Aurora MySQL, multiple cluster database nodes provide increased high availability. There is no single point of failure from the database servers. In addition, since an Aurora cluster can

be distributed across three availability zones, there is a large benefit for high availability and durability of the database. These types of stretch database clusters are usually uncommon with other database architectures.

- Amazon managed storage nodes also provide high availability for the storage tier. A zero-data loss architecture is employed in the event a master node fails and a replica node is promoted to the new master. This failover can usually be performed in under 30 seconds.
- Multiple cluster database nodes enable scaling-out query read performance across multiple servers.
- Greatly reduced operational overhead using a cloud solution and reduced total cost of ownership by using Amazon and open source database engines.
- Automatic management of storage. No need to pre-provision storage for a database. Storage is automatically added as needed, and you only pay for one copy of your data.
- With Amazon Aurora, you can easily scale-out your reads and scale-up your writes which fits perfectly into the workload characteristics of many, if not most, OLTP applications. Scaling out reads usually provides the most tangible performance benefit.

When comparing Oracle RAC and Amazon Aurora side by side, you can see the architectural differences between the two database technologies. Both provide high availability and scalability, but with different architectures.



Overall, Amazon Aurora introduces a simplified solution that can function as an Oracle RAC alternative for many typical OLTP applications that need high performance writes, scalable reads, and very high availability with lower operational overhead.

Feature	Oracle RAC	Amazon Aurora
Storage	Usually enterprise-grade storage + ASM	Aurora Storage Nodes: Distributed, Low Latency, Storage Engine Spanning Multiple AZs
Cluster type	Active/Active. All nodes open for R/W	Active/Active. Primary node open for R/W, Replica nodes open for reads
Cluster virtual IPs	R/W load balancing: SCAN IP	R/W: Cluster endpoint + Read load balancing: Reader endpoint
Internode coordination	Cache-fusion + GCS + GES	N/A
Internode private network	Interconnect	N/A
Transaction (write) TTR from node failure	Typically, 0-30 seconds	Typically, less than 30 seconds
Application (Read) TTR from node failure	Immediate	Immediate
Max number of cluster nodes	Theoretical maximum is 100, but smaller clusters from two to 10 nodes are far more common	15
Provides built-in read scaling	Yes	Yes
Provides built-in write scaling	Yes, under certain scenarios , write performance can be limited and affect scale-out capabilities. For example, when multiple sessions attempt to modify rows	No

Feature	Oracle RAC	Amazon Aurora
	contained in the same database blocks	
Data loss in case of node failure	No data loss	No data loss
Replication latency	N/A	Milliseconds
Operational complexity	Requires database, IT, network, and storage expertise	Provided as a cloud-solution
Scale-up nodes	Difficult with physical hardware, usually requires to replace servers	Easy using the Amazon UI/CLI
Scale-out cluster	Provision, deploy, and configure new servers, unless you pre-allocate a pool of idle servers to scale-out on	Easy using the Amazon UI/CLI

For more information, see [Amazon Aurora as an Alternative to Oracle RAC](#).

Migrate to Aurora MySQL Serverless

Another great option can be Amazon Aurora MySQL using Serverless option, this option it currently available only with Aurora MySQL 5.6 compatible.

Amazon Aurora Serverless is an on-demand, auto-scaling configuration for Amazon Aurora MySQL-compatible edition, where the database will automatically start up, shut down, and scale capacity up or down based on your application's needs. It enables you to run your database in the cloud without managing any database instances. It's a simple, cost-effective option for infrequent, intermittent, or unpredictable workloads.

Manually managing database capacity can take up valuable time and can lead to inefficient use of database resources. With Aurora Serverless, you simply create a database endpoint, optionally specify the desired database capacity range, and connect your applications. You pay on a per-

second basis for the database capacity you use when the database is active, and migrate between standard and serverless configurations with a few clicks in the Amazon RDS Management Console.

For some use cases, this option can be very easy to integrate and it has a big advantage over the Oracle RAC in terms of costs. This instance can be adjusted according to your work load and this is more relevant in terms of scale-out for performance.


You can set the minimum and maximum capacity units required. By doing that, your MySQL serverless instance will scale in/out automatically according to the current workload.

You can choose the following capacity units:

- CPU: 2, RAM: 4 GB
- CPU: 4, RAM: 8 GB
- CPU: 8, RAM: 16 GB
- CPU: 16, RAM: 32 GB
- CPU: 32, RAM: 64 GB
- CPU: 64, RAM: 122 GB
- CPU: 128, RAM: 244 GB
- CPU: 256, RAM: 488 GB

How it works

- Create an Aurora storage volume replicated across multiple AZs.
- Create an endpoint in your VPC for the application to connect to.
- Configure an invisible to the customer network load balancer behind that endpoint.
- Configure multi-tenant request routers to route database traffic to the underlying instances.
- Provision the initial minimum instance capacity.



This option can be easier than using Oracle RAC because with this option you don't need to add or remove servers from the cluster. Also, you don't need to spend on unused hardware, it can scale to even more than you thought you will need when the cluster was created.

For more information, see [Amazon Aurora Serverless](#).

Oracle Traffic Director and Amazon RDS Proxy for Amazon Aurora MySQL

With Amazon DMS, you can migrate Oracle Traffic Director configurations to Amazon Aurora MySQL databases with Amazon RDS Proxy. Oracle Traffic Director is a web server load balancer that distributes client requests across multiple servers. Amazon RDS Proxy for Amazon Aurora MySQL is a fully managed database proxy that facilitates database access with enhanced analytics and performance.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	Some features may be replaced by Amazon RDS Proxy

Oracle usage

Starting with Oracle 18c Oracle Connection Manager can be configured to run in Traffic Director mode. This mode introduces multiple features that help with High Availability, scalability, load balancing, zero downtime and security. Oracle Traffic Director is fast and reliable load-balancing solution. By enabling it for Oracle Connection Manager users can now get following features:

- Increased scalability through usage of transparent connection load-balancing.
- Essential high availability feature of zero downtime that includes support for planned database maintenance, pluggable database relocation, and unplanned database outages for read-mostly workloads.
- High availability of Connection Manager (CMAN) which avoids single point of failure
- Various security features, such as database proxy, firewall, tenant isolation in multi-tenant environment, DDOS protection, and database traffic secure tunneling.

For more information, see [Configuring Oracle Connection Manager in Traffic Director Mode](#) in the *Oracle documentation*.

MySQL usage

Oracle Traffic Director mode for Connection Manager can be potentially replaced by Amazon RDS Proxy for migration to Aurora MySQL.

Amazon RDS Proxy simplifies connection management for Amazon RDS DB instances and clusters. It handles the network traffic between the client application and the database in an active way first by understanding the database protocol. Then Amazon RDS Proxy adjusts its behavior based on the SQL operations from user application and the result sets from the database.

Amazon RDS Proxy also reduces the memory and CPU overhead for the database connection management. The database needs less memory and CPU resources when applications open many simultaneous connections. Amazon RDS Proxy also doesn't require applications to close and reopen connections that stay idle for a long time. Similarly, it requires less application logic to reestablish connections in case of a database problem.

The infrastructure for Amazon RDS Proxy is highly available and deployed over multiple Availability Zones (AZs). The computation, memory, and storage for Amazon RDS Proxy are independent of Amazon RDS DB instances and Aurora DB clusters. This separation helps lower overhead on

database servers, so that they can devote their resources to serving database workloads. The Amazon RDS Proxy compute resources are serverless, automatically scaling based on your database workload.

For more information, see [Amazon RDS Proxy](#) and [Using Amazon RDS Proxy](#) in the *Amazon RDS user guide*.

Oracle Data Pump and MySQL mysqldump and mysql

With Amazon DMS, you can migrate data between different database platforms, including Oracle and MySQL, using tools like Oracle Data Pump and MySQL mysqldump and mysql. Oracle Data Pump provides a streamlined way to move data and metadata from one Oracle database to another, while mysqldump and mysql are utilities for backing up and restoring MySQL databases.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Non-compatible tool.

Oracle usage

Oracle Data Pump is a utility for exporting and importing data from/to an Oracle database. It can be used to copy an entire database, entire schemas, or specific objects in a schema. Oracle Data Pump is commonly used as a part of a backup strategy for restoring individual database objects (specific records, tables, views, stored procedures, and so on) as opposed to snapshots or Oracle RMAN, which provides backup and recovery capabilities at the database level. By default (without using the `sqlfile` parameter during export), the dump file generated by Oracle Data Pump is binary (it can't be opened using a text editor).

Oracle Data Pump supports:

- **Export data from an Oracle database** — The Data Pump EXPDP command creates a binary dump file containing the exported database objects. Objects can be exported with data or with metadata only. Exports can be performed for specific timestamps or Oracle SCNs to ensure cross-object consistency.

- **Import data to an Oracle database** — The Data Pump IMPDP command imports objects and data from a specific dump file created with the EXPDP command. The IMPDP command can filter on import (for example, only import certain objects) and remap object and schema names during import.

The term *logical backup* refers to a dump file created by Oracle Data Pump.

EXPDP and IMPDP can only read or write dump files from file system paths that were pre-configured in the Oracle database as directories. During export or import, users must specify the logical directory name where the dump file should be created; not the actual file system path.

Examples

Use EXPDP to export the HR schema.

```
$ expdp system/**** directory=expdp_dir schemas=hr dumpfile=hr.dmp logfile=hr.log
```

The command contains the credentials to run Data Pump, the logical Oracle directory name for the dump file location (which maps in the database to a physical file system location), the schema name to export, the dump file name, and log file name.

Use IMPDP to import the HR a schema and rename to HR_COPY.

```
$ impdp system/**** directory=expdp_dir schemas=hr dumpfile=hr.dmp logfile=hr.log  
REMAP_SCHEMA=hr:hr_copy
```

The command contains the database credentials to run Data Pump, the logical Oracle directory for where the export dump file is located, the dump file name, the schema to export, the name for the dump file, the log file name, and the REMAP_SCHEMA parameter.

For more information, see [Oracle Data Pump](#) in the *Oracle documentation*.

MySQL usage

MySQL provides two native utilities — mysqldump and mysqlimport. You can use these utilities to perform logical database exports and imports. The functionality is comparable to Oracle's Data Pump utility; however, in some use cases, the mysql connection utility is more equivalent to Oracle Data Pump import tool impdp. You can use these utilities to move data between two different databases or to create logical database backups.

To explain the difference between `mysql` and `mysqlimport` utilities, the equivalent Oracle reference will be used.

Use the `mysql` utility to interact with the database like SQL*Plus. For import purposes, you can run it with all of the `CREATE` and `INSERT` commands to rebuild your schema and insert data just like in Oracle.

The `mysqlimport` utility reads a comma-separated value (CSV) data file and is equivalent to SQL*Loader. This utility is a reference to the `LOAD DATA` command. It is mostly used to move schema or some of the objects between clusters. You should use this utility if you have a data file (not a script) and you want to load it fast.

- `mysqldump` is equivalent to Oracle `expdp`.
- `mysql` is equivalent to Oracle `impdp`.

Amazon Aurora MySQL supports data export and import using `mysqldump`, `mysqlimport`, or `mysql` creation scripts. The binaries for all utilities must be installed on your local workstation or on an Amazon EC2 server.

After export, MySQL dump files created using `mysqldump` can be copied to an Amazon S3 bucket. Later, when the dump files are needed for database restore, they can be copied back to a desktop/server with a MySQL client (such as your workstation or an Amazon EC2 server) to use `mysqlimport`.

- `mysqldump` creates consistent backups only if using the `--single-transaction` option.
- `mysqldump` does not block other readers or writers accessing the database.
- Unlike Data Pump, `mysqldump` files are plain-text.

Examples

Export data using `mysqldump`.

```
mysqldump --column-statistics=0 DATABASE_TO_RESTORE -h INSTANCE_ENDPOINT -P 3306 -u  
USER_NAME -p > /local_path/backup-file.sql
```

Note

In Amazon Relational Database Service (Amazon RDS) for MySQL version 8.0, make sure that the `column_statistics` flag set to 0 if you use binaries when running the `mysqldump`.

Run an export and copy the backup file to an Amazon S3 bucket using a pipe and the Amazon CLI.

```
mysqldump --column-statistics=0 DATABASE_NAME -h MYSQL_INSTANCE_ENDPOINT -P 3306 -u
USER_NAME -p > /local_path/backup-file.sql | aws s3 cp -
s3://mysql-backups/mysql_bck-$(date "+%Y-%m-%d-%H-%M-%S")
```

Import data using `mysql`.

```
mysql DB_NAME -h MYSQL_INSTANCE_ENDPOINT -P 3306 -u
USER_NAME -p < /local_path/backupfile.sql
```

Copy the output file from the local server to an Amazon S3 Bucket using the Amazon CLI.

```
aws s3 cp /local_path/backup-file.sql
s3://my-bucket/backup-$(date "+%Y-%m-%d-%H-%M-%S")
```

In the preceding example, the `{-$(date "+%Y-%m-%d-%H-%M-%S")}` format is valid on Linux servers only.

Download the output file from the S3 bucket.

```
$ aws s3 cp s3://my-bucket/backup-2017-09-10-01-10-10 /local_path/backup-file.sql
```

Summary

Description	Oracle Data Pump	MySQL Dump
Export data to a local file	<pre>expdp system/** schemas=hr dumpfile=hr.dmp logfile=hr.log</pre>	<pre>mysqldump --column- statistics=0 DATABASE_TO_RESTORE -h</pre>

Description	Oracle Data Pump	MySQL Dump
		<pre>INSTANCE_ENDPOINT -P 3306 -u USER_NAME -p > /local_path/backup -file.sql</pre>
Export data to a remote file	<p>Create an Oracle directory on remote storage mount or NFS directory called EXP_DIR. Use the export command:</p> <pre>expdp system/** schemas=hr directory =EXP_DIR dumpfile=hr.dmp logfile=hr.log</pre>	<pre>mysqldump --column- statistics=0 DATABASE_NAME -h MYSQL_INSTANCE_END POINT -P 3306 -u USER_NAME -p > /local_path/backup -file.sql aws s3 cp - s3://mysql- backups/mysql_bck-\$(date +%Y-%m-%d-%H- %M-%S")</pre>
Import data to a new database with a new name	<pre>impdp system/** schemas=hr dumpfile= hr.dmp logfile=hr.log REMAP_SCHEMA=h r:hr_copy TRANSFORMM=OID:N</pre>	<pre>mysql DB_NAME -h MYSQL_INSTANCE_END POINT -P 3306 -u USER_NAME -p < /local_path/backup -file.sql</pre>

For more information, see [mysqldump — A Database Backup Program](#), [mysqlimport — A Data Import Program](#), and [mysql — The MySQL Command-Line Client](#) in the *MySQL documentation*.

Oracle Flashback Database and MySQL snapshots

With Amazon DMS, you can leverage Oracle Flashback Database and MySQL snapshots to restore databases to a previous point in time. Oracle Flashback Database provides a way to rewind a database to a specific time or system change number, while MySQL snapshots capture the state of a database at a particular point in time.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Storage-level backup is managed by Amazon RDS.

Oracle usage

Oracle Flashback Database is a special mechanism built into Oracle databases that helps protect against human errors by providing capabilities to revert the entire database back to a previous point in time using SQL commands. Flashback database implements a self-logging mechanism that captures all changes applied to a database and to data. Essentially, it stores previous versions of database modifications in the configured database Fast Recovery Area.

When you use an Oracle flashback database, you can choose to restore an entire database to either a user-created restore point, a timestamp value, or to a specific System Change Number (SCN).

Examples

Create a database restore point to which you can flashback a database.

```
CREATE RESTORE POINT before_update GUARANTEE FLASHBACK DATABASE;
```

Flashback a database to a previously created restore point.

```
shutdown immediate;
startup mount;
flashback database to restore point before_update;
```

Flashback a database to a specific time.

```
shutdown immediate;
startup mount;
FLASHBACK DATABASE TO TIME "TO_DATE('01/01/2017', 'MM/DD/YY')";
```

For more information, see [FLASHBACK DATABASE](#) in the *Oracle documentation*.

MySQL usage

Snapshots are the primary backup mechanism for Amazon Aurora databases. They are extremely fast and nonintrusive. You can take snapshots using the Amazon RDS Management Console or the Amazon CLI. Unlike RMAN, there is no need for incremental backups. You can choose to restore your database to the exact time when a snapshot was taken or to any other point in time.

Amazon Aurora provides the following types of backups:

- **Automated backups** — Always enabled on Amazon Aurora. They do not impact database performance.
- **Manual backups** — You can create a snapshot at any time. There is no performance impact when taking snapshots of an Aurora database. Restoring data from snapshots requires creation of a new instance. Up to 100 manual snapshots are supported for each database.

When you use Aurora MySQL 5.6 compatible, you can turn on the Aurora Backtrack feature. This feature is equivalent to Flashback Database option in Oracle.

This option applies to newly created MySQL-compatible Aurora database clusters and to MySQL-compatible clusters that have been restored from a backup. Make sure that you opt-in when you create or restore a cluster; you can't turn this option on for a running cluster.

To backtrack your database using Amazon CLI, use the following example. This example demonstrates how to backtrack to instance for one day or 86,400 seconds:

```
aws rds modify-db-cluster
--db-cluster-identifier sample-cluster
--backtrack-window 86400
```

To monitor the backtrack operation, use the following example:

```
aws rds describe-db-cluster-backtracks
--db-cluster-identifier sample-cluster
```

Examples

The following steps to enable Aurora automatic backups and configure the backup retention window as part of the database creation process. This process is equivalent to setting the Oracle

RMAN backup retention policy using the `configure retention policy to recovery window of X days` command.

1. Sign in to your Amazon console and choose **RDS**.
2. Choose **Databases**, then choose your database or create a new one.
3. Expand **Additional configuration** and specify **Backup retention period** in days.

Backup

Backup retention period [Info](#)

Choose the number of days that RDS should retain automatic backups for this instance.


The following table identifies the default automatic backup time for each region.

Region	Default backup window
US West (Oregon)	06:00–14:00 UTC
US West (N. California)	06:00–14:00 UTC
US East (Ohio)	03:00–11:00 UTC
US East (N. Virginia)	03:00–11:00 UTC
Asia Pacific (Mumbai)	16:30–00:30 UTC
Asia Pacific (Seoul)	13:00–21:00 UTC
Asia Pacific (Singapore)	14:00–22:00 UTC
Asia Pacific (Sydney)	12:00–20:00 UTC
Asia Pacific (Tokyo)	13:00–21:00 UTC
Canada (Central)	06:29–14:29 UTC
EU (Frankfurt)	20:00–04:00 UTC

Region	Default backup window
EU (Ireland)	22:00–06:00 UTC
EU (London)	06:00–14:00 UTC
South America (São Paulo)	23:00–07:00 UTC
Amazon GovCloud (US)	03:00–11:00 UTC

Use the following steps to perform a manual snapshot backup of an Aurora database. This process is equivalent to creating a full Oracle RMAN backup (`BACKUP DATABASE PLUS ARCHIVELOG`).

1. Sign in to the Amazon Management Console and choose **RDS**.
2. Choose **Databases**, then choose your database.
3. Choose **Actions** and then choose **Take snapshot**.



Use the following steps to restore an Aurora database from a snapshot. This process is similar to the Oracle RMAN commands `RESTORE DATABASE` and `RECOVER DATABASE`. However, instead of running in place, restoring an Aurora database creates a new cluster.

1. Sign in to the Amazon Management Console and choose **RDS**.
2. Choose **Snapshots**, then choose the snapshot to restore.
3. Choose **Actions** and then choose **Restore snapshot**. This action creates a new instance.

4. On the **Restore snapshot** page, for **DB instance identifier**, enter the name for your restored DB instance.
5. Choose **Restore DB instance**.

Use the following steps to restore an Aurora MySQL database backup to a specific point in time. This process is similar to running the Oracle RMAN command `SET UNTIL TIME 'TO_DATE('XXX')'` before running `RESTORE DATABASE` and `RECOVER DATABASE`.

1. Sign in to the Amazon Management Console and choose **RDS**.
2. Choose **Databases**, then choose your database.
3. Choose **Actions** and then choose **Restore to point in time**.
4. This process launches a new instance. Select the date and time to which you want to restore your database. The selected date and time must be within the configured backup retention for this instance.

Amazon CLI backup and restore operations

In addition to using the Amazon web console to backup and restore an Aurora instance snapshot, you can also use the Amazon CLI to perform the same actions. The CLI is especially useful for migrating existing automated Oracle RMAN scripts to an Amazon environment. The following list highlights some CLI operations:

- Use `describe-db-cluster-snapshots` to view all current Aurora MySQL snapshots.
- Use `create-db-cluster-snapshot` to create a snapshot or restore point.
- Use `restore-db-cluster-from-snapshot` to restore a new cluster from an existing database snapshot.
- Use `create-db-instance` to add new instances to the restored cluster.

```
aws rds describe-db-cluster-snapshots

aws rds create-db-cluster-snapshot
  --db-cluster-snapshot-identifier Snapshot_name
  --db-cluster-identifier Cluster_Name

aws rds restore-db-cluster-from-snapshot
  --db-cluster-identifier NewCluster
```

```

--snapshot-identifier SnapshotToRestore
--engine aurora-mysql

aws rds create-db-instance
--region us-east-1
--db-subnet-group default
--engine aurora-mysql
--db-cluster-identifier NewCluster
--db-instance-identifier newinstance-nodeA
--db-instance-class db.r4.large

```

- Use `restore-db-instance-to-point-in-time` to perform a point-in-time recovery.

```

aws rds restore-db-cluster-to-point-in-time
--db-cluster-identifier clusternamerestore
--source-db-cluster-identifier clustername
--restore-to-time 2017-09-19T23:45:00.000Z

aws rds create-db-instance
--region us-east-1
--db-subnet-group default
--engine aurora-mysql
--db-cluster-identifier clustername-restore
--db-instance-identifier newinstance-nodeA
--db-instance-class db.r4.large

```

Summary

Description	Oracle	Amazon Aurora
Create a restore point	<pre>CREATE RESTORE POINT before_update GUARANTEE FLASHBACK DATABASE;</pre>	<pre>aws rds create-db- cluster-snapshot --db-cluster- snapshotidentifier Snapshot_name --db-cluster-ident ifier Cluster_Name</pre>
Configure flashback retention period	<pre>ALTER SYSTEM SET</pre>	Configure the Backup retention window

Description	Oracle	Amazon Aurora
	<pre>db_flashback_reten tion_target=2880;</pre>	<p>setting using the Amazon Management Console or Amazon CLI.</p>
<p>Flashback database to a previous restore point</p>	<pre>shutdown immediate; startup mount; flashback database to restore point before_update;</pre>	<p>Create a new cluster from a snapshot.</p> <pre>aws rds restore-db- cluster-from-snapshot --db-cluster-ident ifier NewCluster --snapshot-identif ier SnapshotToRestore --engine aurora-my sql</pre> <p>Add a new instance to the cluster.</p> <pre>aws rds create-db- instance --region useast-1 --db-subnet-group default --engine aurora-my sql --db-cluster-ident ifier clustername- restore --db-instanceident ifier newinstance- nodeA --dbinstance-class db.r4.large</pre>

Description	Oracle	Amazon Aurora
Flashback database to a previous point in time	<pre>shutdown immediate; startup mount; FLASHBACK DATABASE TO TIME "TO_DATE ('01/01/2 017', 'MM/DD/YY')";</pre>	<p>Use the following example to restore your database to 86,400 seconds ago.</p> <pre>aws rds modify-db- cluster --db-clusteridenti fier sample-cluster --backtrack-window 86400</pre>

For more information, see [mysqldump — A Database Backup Program](#) in the *MySQL documentation*, [rds](#) in the *CLI Command Reference* and [Restoring a DB instance to a specified time](#) and [Restoring from a DB snapshot](#) in the *Amazon RDS user guide*.

Oracle Flashback Table and MySQL snapshots

With Amazon DMS, you can restore databases to a specific point in time using Oracle Flashback Table and MySQL snapshots. Oracle Flashback Table provides a way to view and restore data from a specified time in the past, while MySQL snapshots capture the state of a database at a specific point for backup or replication purposes.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Storage-level backup managed by Amazon RDS.

Oracle usage

Oracle Flashback Table is a data protection feature used to undo changes to a table and rewind it to a previous state, not from the backup. While Flashback table operations are running, the affected tables are locked, but the rest of the database remains available.

If the structure of a table has been changed since the point of restore, the FLASHBACK will fail.

Make sure that the row movement is turned on.

The data to restore must be found in the undo, and the database administrator manages the size and retention.

You can restore a table to a System Change Number (SCN), restore point, or timestamp.

Examples

Flashback a table using SCN (query V\$DATABASE to obtain the SCN).

```
SELECT CURRENT_SCN FROM V$DATABASE;  
FLASHBACK TABLE employees TO SCN 3254648;
```

Flashback a table using a restore point (query V\$RESTORE_POINT to obtain restore points).

```
SELECT NAME, SCN, TIME FROM V$RESTORE_POINT;  
FLASHBACK TABLE employees TO RESTORE POINT employees_year_update;
```

Flashback a table using a timestamp (query V\$PARAMETER to obtain the undo_retention value).

```
SELECT NAME, VALUE/60 MINUTES_RETAINED  
FROM V$PARAMETER  
WHERE NAME = 'undo_retention';  
FLASHBACK TABLE employees TO  
TIMESTAMP TO_TIMESTAMP('2017-09-21 09:30:00', 'YYYY-MM-DD HH:MI:SS');
```

For more information, see [Backup and Recovery User Guide](#) in the *Oracle documentation*.

MySQL usage

Snapshots are the primary backup mechanism for Amazon Aurora databases. They are extremely fast and nonintrusive. You can take snapshots using the Amazon RDS Management Console or the

Amazon CLI. Unlike RMAN, there is no need for incremental backups. You can choose to restore your database to the exact time when a snapshot was taken or to any other point in time.

Amazon Aurora provides the following types of backups:

- **Automated backups** — Always enabled on Amazon Aurora. They do not impact database performance.
- **Manual backups** — You can create a snapshot at any time. There is no performance impact when taking snapshots of an Aurora database. Restoring data from snapshots requires creation of a new instance. Up to 100 manual snapshots are supported for each database.

Examples

For examples, see [MySQL Snapshots](#).

Summary

Description	Oracle	Amazon Aurora
Create a restore point	<pre>CREATE RESTORE POINT before_update GUARANTEE FLASHBACK DATABASE;</pre>	<pre>aws ids create-db- cluster-snapshot --db-cluster- snapshotidentifier Snapshot_name --db-cluster-ident ifier Cluster_Name</pre>
Configure flashback retention period	<pre>ALTER SYSTEM SET db_flashback_reten tion_target=2880;</pre>	Configure the Backup retention window setting using the Amazon Management Console or Amazon CLI.
Flashback table to a previous restore point	<pre>shutdown immediate; startup mount; flashback database to restore point before_update;</pre>	Create new cluster from a snapshot. <pre>aws ids restore-db- cluster-from-snapshot</pre>

Description	Oracle	Amazon Aurora
		<pre data-bbox="1073 205 1503 468">--db-cluster-identifier NewCluster --snapshot-identifier SnapshotToRestore --engine aurora-mysql</pre> <p data-bbox="1073 499 1503 583">Add new instance to the cluster.</p> <pre data-bbox="1073 615 1503 1255">aws rds create-db-instance --region useast-1 --db-subnet-group default --engine aurora-mysql --db-cluster-identifier clustername-restore --db-instanceidentifier newinstance-nodeA --dbinstance-class db.r4.large</pre> <p data-bbox="1073 1287 1503 1476">Use <code>mysqldbexport</code> and <code>mysql</code> to copy the table from the restored instance to the original instance.</p>

Description	Oracle	Amazon Aurora
Flashback table to a previous point in time	<pre>shutdown immediate; startup mount; FLASHBACK DATABASE TO TIME "TO_DATE ('01/01/2 017', 'MM/DD/YY')";</pre>	<p>Create a new cluster from a snapshot and provide a specific point in time.</p> <pre>aws rds restore-db- cluster-to-point-in- time --db-cluster-ident ifier clustername- restore --source-db-cluster- identifier clustername --restore-to-time 2017-09-19T23:45:0 0.000Z</pre> <p>Add a new instance to the cluster:</p> <pre>aws rds create-db- instance --region us-east-1 --db-subnetgroup default --engine aurora-mysql --db-cluster-ident ifier clustername- restore --db-instance-iden tifier newinstance- nodeA --db-instance-class db.r4.large</pre> <p>Use <code>mysqldbexport</code> and <code>mysql</code> to copy the table from the restored instance to the original instance.</p>

For more information, see [mysqldump — A Database Backup Program](#) in the *MySQL documentation*, [rds](#) in the *CLI Command Reference* and [Restoring a DB instance to a specified time](#) and [Restoring from a DB snapshot](#) in the *Amazon RDS user guide*.

Oracle Recovery Manager and Amazon RDS snapshots

With Amazon DMS, you can migrate data from Oracle databases by using Oracle Recovery Manager (RMAN) backup sets or Amazon RDS snapshots. Oracle Recovery Manager is a utility for backing up, restoring, and recovering Oracle databases. Amazon RDS snapshots capture the entire database instance, including transaction logs, at a specific point in time.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Storage-level backup managed by Amazon RDS.

Oracle usage

Oracle Recovery Manager (RMAN) is a primary backup and recovery tool in Oracle. It provides its own scripting syntax and can be used to take full or incremental backups of an Oracle database. The following list identifies the types of backups.

- **Full RMAN backup** — Creates a full backup of an entire database or individual Oracle data files. For example, a level 0 full backup.
- **Differential incremental RMAN backup** — Performs a backup of all database blocks that have changed from the previous level 0 or 1 backup.
- **Cumulative incremental RMAN backup** — Perform a backup all of blocks that have changed from the previous level 0 backup.

RMAN supports online backups of an Oracle database if it has been configured to run in Archived Log Mode.

RMAN backs up the following files:

- Database data files.
- Database control file.
- Database parameter file.
- Database Archived Redo Logs.

Examples

Use the RMAN CLI to connect to an Oracle database.

```
export ORACLE_SID=ORCL
rman target=/
```

Perform a full backup of the database and the database archived redo logs.

```
BACKUP DATABASE PLUS ARCHIVELOG;
```

Perform an incremental level 0 or level 1 backup of the database.

```
BACKUP INCREMENTAL LEVEL 0 DATABASE;
BACKUP INCREMENTAL LEVEL 1 DATABASE;
```

Restore a database.

```
RUN {
SHUTDOWN IMMEDIATE;
STARTUP MOUNT;
RESTORE DATABASE;
RECOVER DATABASE;
ALTER DATABASE OPEN;
}
```

Restore a specific pluggable database (Oracle 12c).

```
RUN {
ALTER PLUGGABLE DATABASE pdbA, pdbB CLOSE;
RESTORE PLUGGABLE DATABASE pdbA, pdbB;
RECOVER PLUGGABLE DATABASE pdbA, pdbB;
ALTER PLUGGABLE DATABASE pdbA, pdbB OPEN;
}
```

Restore a database to a specific point in time.

```
RUN {  
SHUTDOWN IMMEDIATE;  
STARTUP MOUNT;  
SET UNTIL TIME "TO_DATE('20-SEP-2017 21:30:00', 'DD-MON-YYYY HH24:MI:SS')";  
RESTORE DATABASE;  
RECOVER DATABASE;  
ALTER DATABASE OPEN RESETLOGS;  
}
```

List all current database backups created with RMAN.

```
LIST BACKUP OF DATABASE;
```

For more information, see [Backup and Recovery User Guide](#) in the *Oracle documentation*.

MySQL usage

Snapshots are the primary backup mechanism for Amazon Aurora databases. They are extremely fast and nonintrusive. You can take snapshots using the Amazon RDS Management Console or the Amazon CLI. Unlike RMAN, there is no need for incremental backups. You can choose to restore your database to the exact time when a snapshot was taken or to any other point in time.

Amazon Aurora provides the following types of backups:

- **Automated backups** — Always enabled on Amazon Aurora. They do not impact database performance.
- **Manual backups** — You can create a snapshot at any time. There is no performance impact when taking snapshots of an Aurora database. Restoring data from snapshots requires creation of a new instance. Up to 100 manual snapshots are supported for each database.

Note

In Amazon Relational Database Service (Amazon RDS) for MySQL version 8.0.21, you can turn on or off the redo logging option using the `ALTER INSTANCE {ENABLE|DISABLE} INNODB REDO_LOG` syntax. This functionality is intended for loading data into a new MySQL instance. Turning off the redo logging option helps speed up data loading by avoiding redo log writes. The new `INNODB_REDO_LOG_ENABLE` privilege permits turning

on and turning off the redo logging option. The new `Innodb_redo_log_enabled` status variable permits monitoring redo logging status. For more information, see [Disabling Redo Logging](#) in the *MySQL documentation*.

Examples

For examples, see [MySQL Snapshots](#).

Summary

Description	Oracle	Amazon Aurora
Scheduled backups	Create the <code>DBMS_SCHEDULER</code> job that will run your RMAN script on a scheduled basis.	Automatic
Manual full database backups	<pre>BACKUP DATABASE PLUS ARCHIVELOG;</pre>	<p>Use Amazon RDS dashboard or the Amazon CLI command to take a snapshot on the cluster.</p> <pre>aws rds create-db-cluster-snapshot --dbcluster-snapsh ot-identifier Snapshot_name --db-cluster-ident ifier Cluster_Name</pre>
Restore database	<pre>RUN { SHUTDOWN IMMEDIATE; STARTUP MOUNT; RESTORE DATABASE; RECOVER DATABASE; ALTER DATABASE OPEN; }</pre>	<p>Create new cluster from a cluster snapshot.</p> <pre>aws rds restore-db-cluster-from-snapshot --db-cluster-ident ifier NewCluster</pre>

Description	Oracle	Amazon Aurora
		<pre data-bbox="1073 212 1507 344">--snapshotidentifier SnapshotToRestore --engine aurora-mysql</pre> <p data-bbox="1068 386 1446 464">Add a new instance to the new/restored cluster.</p> <pre data-bbox="1073 506 1507 1100">aws rds create-db- instance --region useast-1 --db-subnet-group default --engine aurora-mysql --db-cluster-ident ifier clustername- restore --db-instance-iden tifier newinstance- nodeA --db-instance-class db.r4.large</pre>
Incremental differential	<pre data-bbox="597 1136 1019 1415">BACKUP INCREMENTAL LEVEL 0 DATABASE; BACKUP INCREMENTAL LEVEL 1 DATABASE;</pre>	N/A
Incremental cumulative	<pre data-bbox="597 1451 1019 1730">BACKUP INCREMENTAL LEVEL 0 CUMULATIVE DATABASE; BACKUP INCREMENTAL LEVEL 1 CUMULATIVE DATABASE;</pre>	N/A

Description	Oracle	Amazon Aurora
Restore a database to a specific point in time	<pre> RUN { SHUTDOWN IMMEDIATE; STARTUP MOUNT; SET UNTIL TIME "TO_DATE('19-SEP-2017 23:45:00', 'DD-MON-YYYY HH24:MI:SS')"; RESTORE DATABASE; RECOVER DATABASE; ALTER DATABASE OPEN RESETLOGS; } </pre>	<p>Create a new cluster from a cluster snapshot by given custom time to restore.</p> <pre> aws rds restore-db-cluster-to-point-in-time --db-cluster-identifier clustername-restore --source-db-cluster-identifier clustername --restore-to-time 2017-09-19T23:45:00.000Z </pre> <p>Add a new instance to the new or restored cluster.</p> <pre> aws rds create-db-instance --region useast-1 --db-subnet-group default --engine aurora-mysql --db-cluster-identifier clustername-restore --db-instance-identifier newinstance-nodeA --db-instance-class db.r4.large </pre>
Backup database archive logs	<pre> BACKUP ARCHIVELOG ALL; </pre>	N/A

Description	Oracle	Amazon Aurora
Delete old database archive logs	<pre>CROSSCHECK BACKUP; DELETE EXPIRED BACKUP;</pre>	N/A

Description	Oracle	Amazon Aurora
Restore a single pluggable database (12c)	<pre> RUN { ALTER PLUGGABLE DATABASE pdb1, pdb2 CLOSE; RESTORE PLUGGABLE DATABASE pdb1, pdb2; RECOVER PLUGGABLE DATABASE pdb1, pdb2; ALTER PLUGGABLE DATABASE pdb1, pdb2 OPEN; } </pre>	<p>Create new cluster from a cluster snapshot.</p> <pre> aws rds restore-db- cluster-from-snapshot --db-cluster-ident ifier NewCluster --snapshotidentifier SnapshotToRestore --engine aurora-mysql </pre> <p>Add a new instance to the new or restored cluster.</p> <pre> aws rds create-db- instance --region useast-1 --db-subnet-group default --engine aurora-mysql --db-cluster-ident ifier clustername- restore --db-instance-iden tifier newinstance- nodeA --db-instance-class db.r4.large </pre> <p>Use mysqldump and mysql to copy the database to the original instance.</p> <pre> mysqldump --column- statistics=0 DATABASE_TO_RESTOR E -h RESTORED_ INSTANCE_ENDPOINT </pre>

Description	Oracle	Amazon Aurora
		<pre data-bbox="1073 212 1503 386">-P 3306 -u USER_NAME -p > /local_path/ backup-file.sql</pre> <pre data-bbox="1073 422 1503 730">mysql DB_NAME -h MYSQL_INS TANCE_ENDPOINT -P 3306 -u USER_NAME -p < /local_path/ backup-file.sql</pre> <div data-bbox="1073 772 1503 1276"> <p>Note</p> <p>In Amazon RDS for MySQL version 8.0, make sure that the <code>column_statistics</code> flag is set to 0 if you use binaries when running the <code>mysqldump</code>.</p> </div>

For more information, see [mysqldump — A Database Backup Program](#) in the *MySQL documentation*, [rds](#) in the *CLI Command Reference*, [Restoring a DB instance to a specified time](#) and [Restoring from a DB snapshot](#) in the *Amazon RDS user guide*.

Oracle SQL*Loader and MySQL mysqlimport and LOAD DATA

With Amazon DMS, you can efficiently migrate data from flat files into Amazon databases using Oracle SQL*Loader, MySQL `mysqlimport`, and `LOAD DATA` commands. These utilities facilitate bulk data loading from external files into database tables.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	The tool isn't compatible.

Oracle usage

SQL*Loader is a powerful utility that imports data from external files into database tables. It has strong parsing engine with few limitations on data formats.

You can use SQL*Loader with or without a control file. A control file enables handling more complicated load environments. For simpler loads, use SQL*Loader without a control file. The same also refers to SQL*Loader Express.

The outputs of SQL*Loader include the imported database data, a log file, a bad file or rejected records, and a discard file, if this option is turned on.

Examples

Oracle SQL*Loader is well suited for large databases with a limited number of objects. The process of exporting from a source database and loading to a target database is very specific to the schema. The following example creates sample schema objects, exports from a source, and loads into a target database.

Create a source table.

```
CREATE TABLE customer_0 TABLESPACE users
AS SELECT rownum id, o.* FROM all_objects o, all_objects x
where rownum <= 1000000;
```

On the target Amazon RDS instance, create a destination table for the loaded data.

```
CREATE TABLE customer_1 TABLESPACE users
AS select 0 as id, owner, object_name, created
from all_objects where 1=2;
```

The data is exported from the source database to a flat file with delimiters. This example uses SQL*Plus. For your data, you will likely need to generate a script that does the export for all the objects in the database.

```
alter session set nls_date_format = 'YYYY/MM/DD HH24:MI:SS';
set linesize 800
HEADING OFF FEEDBACK OFF array 5000 pagesize 0
spool customer_0.out
SET MARKUP HTML PREFORMAT ON SET COLSEP ',' SELECT id,
  owner, object_name, created FROM customer_0;
spool off
```

Create a control file describing data. Depending on data, you may need to build a script that provides this functionality.

```
cat << EOF > sqlldr_1.ctl
LOAD DATA
INFILE customer_0.out
into table customer_1
APPEND
fields terminated by "," optionally enclosed by '"'
(id POSITION(01:10) INTEGER EXTERNAL,
owner POSITION(12:41) CHAR,
object_name POSITION(43:72) CHAR,
created POSITION(74:92) date "YYYY/MM/DD HH24:MI:SS")
```

Import data using SQL*Loader. Use the appropriate user name and password for the target database.

```
sqlldr cust_dba@targetdb control=sqlldr_1.ctl BINDSIZE=10485760 READSIZE=10485760
ROWSS=1000
```

For more information, see [SQL*Loader](#) in the *Oracle documentation*.

MySQL usage

You can use the two following options as a replacement for the Oracle SQL*Loader utility:

- **MySQL Import** using an export file similar to a control file.
- **Load from Amazon S3 File** using a table-formatted file on Amazon S3 and loading it into a MySQL database.

MySQL Import is a good option when you can use a tool from another server or a client. The LOAD DATA command can be combined with metadata tables and EVENT objects to schedule loads.

For more information, see [Loading data into an Amazon Aurora MySQL DB cluster from text files in an Amazon S3 bucket](#) in the *User Guide for Aurora* and [mysqlimport — A Data Import Program](#) in the *MySQL documentation*.

Oracle and MySQL configuration

This section provides pages about Oracle and MySQL configuration topics.

Topics

- [Oracle and Aurora MySQL upgrades](#)
- [Oracle alert log and MySQL error log](#)
- [Oracle SGA and PGA memory sizing and MySQL memory buffers](#)
- [Oracle instance parameters and Aurora MySQL parameter groups](#)
- [Oracle session parameters and MySQL session variables](#)

Oracle and Aurora MySQL upgrades

With Amazon DMS, you can upgrade your Oracle and Aurora MySQL databases to newer versions with minimal downtime. The Oracle and Aurora MySQL Upgrades feature facilitates seamless database upgrades by creating a new database instance with the desired version, migrating data from the old instance, and redirecting applications to the new instance. This capability is crucial for organizations that need to stay current with the latest database software releases for security, performance, and compatibility reasons.

Oracle usage

As a Database Administrator, from time to time a database upgrade is required, it can be either for security fix, but, or a new database feature.

The Oracle upgrades are divided into two different types of upgrades, minor and major.

This topic will outline the differences between the procedure to execute upgrades on your Oracle databases today and how you will run those upgrades post migrating to Amazon RDS running Aurora.

The regular presentation of Oracle versions is combined of 4 numbers divided by dots, and sometimes you can see the fifth number.

Either way, major or minor upgrades, the first step to initiate the processes mentioned above would be to install the new Oracle software on the database server, and of course before upgrading

a production database to have an extensive amount of testing with the applications using the database to upgrade.

Oracle 18c introduces Zero-Downtime Database Upgrade to automate Database upgrade and potentially eliminate application downtime during this process.

To understand the versions, let us use the following example 11.2.0.4.0. The digits in this example mean the following:

- 11 — The major database version.
- 2 — The database maintenance version.
- 0 — The application server version.
- 4 — The component-specific version.
- 0 — The platform-specific version.

For more information, see [About Oracle Database Release Numbers](#) in the *Oracle documentation*.

In Oracle, users can set the compatibility level of the database to control the features and some behaviors.

You can do this by using the COMPATIBLE parameter. Use the following query to fetch the value for this parameter:

```
SELECT NAME, VALUE FROM V$PARAMETER WHERE NAME = 'compatible';
```

Upgrade process

In general, the process of major or minor upgrades is the same. Minor version upgrade has less steps but overall the process is similar.

Major upgrade referring to upgrades of the version number in the Oracle version. In the preceding example, the minor upgrade refers to any of the following numbers in the Oracle version that follow the major database version: 2.0.4.0.

Major upgrades are mostly being done in order to gain many new useful features being released between those versions, while minor upgrades are focused on bug and security fixes.

You can upgrade your database version using the Oracle upgrade tools or manually.

Oracle tools perform the following steps and might ask for some inputs or fixes from the user along the process. The upgrade steps are:

- **Upgrade operation type** — the user chooses either to upgrade an Oracle database or move a database between Oracle software installations.
- **Database selection** — the user selects the database to upgrade and the Oracle software to use for this database.
- **Prerequisite checks** — Oracle tools let the user choose what to do with all issues found and their severity.
- **Upgrade options** — Oracle lets the user to pick his practices to do the upgrade, using such options as: recompilation and parallelism for those, time zone upgrade, statistics gathering, and more.
- **Management options** — the user chooses to connect or configure the Oracle management solutions to the database.
- **Move database files** — the user chooses if a data file movement is required to a new devices or path.
- **Network configuration** — Oracle listener configurations.
- **Recovery options** — the user defines Oracle backup solutions or uses his own.
- **Summary** — a report of all options were selected in previous steps to present before the upgrade.
- **Progress** — monitor and present the upgrade status.
- **Results** — a post-upgrade summary.

For the manual process, we won't cover all actions in this topic, as there are many steps and commands to run.

In overall, the eleven steps mentioned before, will be divided into many sub-steps and tasks to run.

For more information, see [Example of Manual Upgrade of Windows Non-CDB Oracle Database 11.2.0.3](#) in the *Oracle documentation*.

MySQL usage

After migrating your databases to Amazon Aurora MySQL, make sure that you upgrade your database instance from time to time, for the same reasons you have done it in the past like new features, bugs and security fixes.

In a managed service such as Amazon Relational Database Service, the upgrade process is much easier and simpler compare to the on-prem Oracle process.

To determine the current Aurora for MySQL version being used, you can use the following Amazon CLI command:

```
aws rds describe-db-engine-versions --engine aurora-mysql --query '*[].[EngineVersion]' --output text --region your-AWS-Region
```

You can also query this from the database, using the following queries:

```
SELECT AURORA_VERSION();
```

In an Aurora MySQL version number scheme, for example 2.08.1, first digit represents the major version. Aurora MySQL version 1 is compatible with MySQL 5.6 and Aurora MySQL version 2 is compatible with MySQL 5.7. For more information about Aurora and MySQL versions mapping, see [Database engine updates for Amazon Aurora MySQL version 2](#) in the *Release Notes for Aurora MySQL*.

Amazon doesn't apply major version upgrades on Amazon RDS Aurora automatically. Major version upgrades contains new features and functionality which often involves system table and other code changes. These changes may not be backward-compatible with previous versions of the database so applications testing is highly recommended.

Applying automatic minor upgrades can be set by configuring the Amazon RDS instance to allow it.

You can use the following Amazon CLI command on Linux to determine the current automatic upgrade minor versions.

```
aws rds describe-db-engine-versions --output=table --engine mysql --engine-version minor-version --region region
```

Note

If the query doesn't return results, there is no automatic minor version upgrade available and scheduled.

When enabled, the instance will be automatically upgraded during the scheduled maintenance window.

To upgrade your cluster to a compatible cluster, you can do so by running an upgrade process on the cluster itself. This kind of upgrade is an in-place upgrade, in contrast to upgrades that you do by creating a new cluster. The upgrade is relatively fast because it doesn't require copying all your data to a new cluster volume. In-place upgrade preserves the endpoints and set of DB instances for your cluster.

To verify application compatibility, performance and maintenance procedures for the upgraded cluster, you can perform a simulation of the upgrade by doing the following: * Clone a cluster. * Perform an in-place upgrade of the cloned cluster. * Test applications, performance and so on, using the cloned cluster. * Resolve any issues, adjust your upgrade plans to account for them. * Once all the testing looks good, you can perform the in-place upgrade for your production cluster.

For major upgrades, Amazon recommends the following flow:

- Check for open XA transactions by running the `XA RECOVER` statement. Commit or rollback the XA transactions before starting the upgrade.
- Check for DDL statements by running the `SHOW PROCESSLIST` statement and looking for `CREATE`, `DROP`, `ALTER`, `RENAME`, and `TRUNCATE` statements in the output. Allow all DDLs to finish before starting the upgrade.
- Check for any uncommitted rows by querying the `INFORMATION_SCHEMA.INNODB_TRX` table. The table contains one row for each transaction. Let the transaction complete or shut down applications that are submitting these changes.

Aurora MySQL performs a major version upgrade in multiple steps. After each step begins, Aurora MySQL records an event. You can monitor the current status and events as they occur on the Events page in the Amazon RDS console.

Aurora performs a series of checks before beginning the upgrade process. If any issues are detected during these checks, resolve the issue identified in the event details and restart the upgrade process.

Aurora takes the cluster offline, performs a similar set of tests as in the previous step. If no new issues are identified, then Aurora moves with the next step. If any issues are detected during these checks, resolve the issue identified in the event details and restart the upgrade process again.

Aurora backs up the MySQL cluster by creating a snapshot of the cluster volume.

Aurora clones the cluster volume. If any issues are encountered during the upgrade, Aurora reverts to the original data from the cloned cluster volume and brings the cluster back online.

Aurora performs a clean shutdown and it rolls back any uncommitted transactions.

Aurora upgrades the engine version. It installs the binary for the new engine version and uses the writer DB instance to upgrade your data to new to MySQL compatible format. During this stage, Aurora modifies the system tables and performs other conversions that affect the data in your cluster volume.

The upgrade process is completed. Aurora records a final event to indicate that the upgrade process completed successfully. Now DB cluster is running the new major version.

Upgrade can be done through the Amazon Console or Amazon CLI.

Upgrade using the Amazon Management Console

1. Sign in to the Amazon Management Console and choose **RDS**.
2. Choose **Databases**, and then choose the DB cluster that you want to upgrade.
3. Choose **Modify**. The **Modify DB cluster page** appears.
4. For **DB engine version**, choose the new version.
5. Choose **Continue** and check the summary of modifications.
6. To apply the changes immediately, choose **Apply immediately**. Choosing this option can cause an outage in some cases. For more information, see [Modifying an Amazon Aurora DB cluster](#) in the *User Guide for Aurora*.
7. On the confirmation page, review your changes. If they are correct, choose **Modify cluster** to save your changes. Alternatively, choose **Back** to edit your changes or **Cancel** to cancel your changes.

Upgrade using Amazon CLI

To upgrade the major version of an Aurora MySQL DB cluster, use the Amazon CLI `modify-db-cluster` command with the following required parameters:

For Linux, macOS, or Unix:

```
aws rds modify-db-cluster \  
--db-cluster-identifier sample-cluster \  
--engine aurora-mysql \  
--engine-version 5.7.mysql_aurora.2.09.0 \  

```

```
--allow-major-version-upgrade \  
--apply-immediately
```

For Windows:

```
aws rds modify-db-cluster ^  
--db-cluster-identifier sample-cluster ^  
--engine aurora-mysql ^  
--engine-version 5.7.mysql_aurora.2.09.0 ^  
--allow-major-version-upgrade ^  
--apply-immediately
```

Summary

Phase	Oracle	Aurora MySQL
Prerequisite	Install new Oracle software	N/A
Prerequisite	Upgrade operation type	N/A
Prerequisite	Database selection	Select the right Amazon Relational Database Service instance
Prerequisite	Prerequisite checks	Commit or rollback uncommitted transactions
Prerequisite	Upgrade options	N/A
Prerequisite	Management options (optional)	N/A
Prerequisite	Move database files (optional)	N/A
Prerequisite	Network configuration (optional)	N/A
Prerequisite	Recovery options	N/A
Prerequisite	Summary	N/A

Phase	Oracle	Aurora MySQL
Prerequisite	Perform a database backup	Run an Amazon Relational Database Service instance backup
Prerequisite	Stop application and connection	Stop application and connection
Run	Progress	Can be reviewed from the console
Post-upgrade	Results	Can be reviewed from the console
Post-upgrade	Test applications against the new upgraded database	Test applications against the new upgraded database
Production deployment	Re-run all steps in a production environment	Re-run all steps in a production environment

For more information, see [Upgrading Amazon Aurora MySQL DB clusters](#) in the *User Guide for Aurora*.

Oracle alert log and MySQL error log

With Amazon DMS, you can capture and analyze Oracle alert log and MySQL error log during database migration tasks. The Oracle alert log records notifications and errors raised by the Oracle database, while the MySQL error log tracks errors, warnings, and notes generated by the MySQL server.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Use Event Notifications Subscription

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
			with Amazon Simple Notification Service (SNS) .

Oracle usage

The primary Oracle error log file is the *alert log*. It contains verbose information about database activity including informational messages and errors. Each event includes a timestamp indicating when the event occurred. The alert log filename format is `alert<sid>.log`.

The alert log is the first place to look when troubleshooting or investigating errors, failures, and other messages indicating a potential database problem. Common events logged in the alert log include:

- Database startup or shutdown.
- Database redo log switch.
- Database errors and warnings, which begin with ORA- followed by an Oracle error number.
- Network and connection issues.
- Links for a detailed trace files about specific database events.

The Oracle Alert Log can be found inside the database Automatic Diagnostics Repository (ADR), which is a hierarchical file-based repository for diagnostic information: `$ADR_BASE/diag/rdbms/{DB-name}/{SID}/trace`.

In addition, several other Oracle server components have unique log files such as the database listener and the Automatic Storage Manager (ASM).

Examples

The following screenshot displays partial contents of the Oracle database alert log file.

```

Sun Sep 03 13:27:23 2017
Starting ORACLE instance (normal)
***** Large Pages Information *****
Per process system memlock (soft) limit = 64 KB

Total Shared Global Region in Large Pages = 0 KB (0%)

Large Pages used by this instance: 0 (0 KB)
Large Pages unused system wide = 0 (0 KB)
Large Pages configured system wide = 0 (0 KB)
Large Page size = 2048 KB

```

For more information, see [Monitoring Errors and Alerts](#) in the *Oracle documentation*.

MySQL usage

MySQL provides detailed logging and reporting of errors that occur during the database and connected sessions life cycle. In an Amazon Aurora deployment, these informational and error messages are accessible using the Amazon Relational Database Service console.

MySQL and Oracle error codes

Oracle	MySQL
ORA-00001: unique constraint string.st ring violated.	Error [1062][23000]: Duplicate entry value for key column.

For more information, see [Server Error Message Reference](#) in the *MySQL documentation*.

Error log types

MySQL provides several types of logs.

Log type	Information written to log
Error log	Problems encountered starting, running, or stopping mysqld.
General query log	Established client connections and statements received from clients.

Log type	Information written to log
Binary log	Statements that change data (also used for replication).
Relay log	Data changes received from a replication master server.
Slow query log	Queries that took more than <code>long_query_time</code> seconds to execute.
DDL log (metadata log)	Meta-data operations performed by DDL statements.

For more information, see [MySQL Server Logs](#) in the *MySQL documentation*.

Examples

Access the MySQL error log using the Amazon Relational Database Service or Amazon Aurora console:

1. Sign in to the Amazon Management Console, choose **RDS**, and then choose **Databases**.
2. Choose the instance name.
3. Choose **Logs & events** and select the log to inspect. For example, select the log during the hour the data was experiencing problems. The following screenshot displays partial contents of a MySQL database error log as viewed from the Amazon Relational Database Service console.

Viewing Log: error/mysql-error-running.log.2018-02-09.19 (93 kB)

text: ■ background: ■

```

2018-02-09T19:01:48.847682Z 0 [Warning] 'NO_AUTO_CREATE_USER' sql mode was not set.
2018-02-09T19:01:48.848471Z 0 [Warning] Insecure configuration for --secure-file-priv:
Location is accessible to all OS users. Consider choosing a different directory.
2018-02-09T19:01:48.848497Z 0 [Note] /rdsdbbin/oscar/bin/mysqld (mysqld 5.7.12-log)
starting as process 5478 ...
2018-02-09T19:01:48.881215Z 0 [Warning] InnoDB: Setting innodb_checksums to OFF is
DEPRECATED. This option may be removed in future releases. You should set
innodb_checksum_algorithm=NONE instead.
2018-02-09T19:01:48.881317Z 0 [Note] InnoDB: PUNCH HOLE support not available
2018-02-09T19:01:48.881329Z 0 [Note] InnoDB: Mutexes and rw_locks use GCC atomic
builtins
2018-02-09T19:01:48.881333Z 0 [Note] InnoDB: Uses event mutexes
2018-02-09T19:01:48.881336Z 0 [Note] InnoDB: GCC builtin __atomic_thread_fence() is
used for memory barrier
2018-02-09T19:01:48.881340Z 0 [Note] InnoDB: Compressed tables use zlib 1.2.3
2018-02-09T19:01:48.882024Z 0 [Note] InnoDB: == Add SYNC_FAST DDL..
2018-02-09T19:01:48.884857Z 0 [Note] InnoDB: Number of pools: 1
2018-02-09T19:01:48.888837Z 0 [Note] InnoDB: Using CPU crc32 instructions
2018-02-09T19:01:48.890231Z 0 [Note] InnoDB: Initializing buffer pool, total size =

```

MySQL error log configuration

Several parameters specify the location of the MySQL log and errors files. The following table identifies common Amazon Aurora configuration options.

Parameter	Description
log_error	Sets the file name and path for the error log. You can modify it through an Aurora Database Parameter Group.
log_error_verbosity	Sets the message levels that are logged such as error, warning, note messages, and so on. You can modify it through an Aurora Database Parameter Group.

Parameter	Description
USE_SLOW_LOG	Sets the minimum execution time above which statements are logged in ms. You can modify it through an Aurora Database Parameter Group.

Note

Modifications of certain parameters, such as `log_error` are turned off for Aurora MySQL instances.

Oracle SGA and PGA memory sizing and MySQL memory buffers

With Amazon DMS, you can optimize memory utilization for migrated databases by configuring Oracle System Global Area (SGA) and Program Global Area (PGA) memory sizes, as well as MySQL memory buffers. The SGA is a group of shared memory structures that contain data and control information for an Oracle database instance, while the PGA is a memory region that contains data and control information for a server process. MySQL memory buffers, such as the buffer pool and query cache, manage data caching and query performance.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Different cache names, similar usage.

Oracle usage

An Oracle instance allocates several individual pools of server RAM used as various caches for the database. These include the Buffer Cache, Redo Buffer, Java Pool, Shared Pool, Large Pool, and others. The caches reside in the System Global Area (SGA) and are shared across all Oracle sessions.

In addition to the SGA, each Oracle session is granted an additional area of memory for session-private operations (sorting, private SQL cursors elements, and so on) called the Private Global Area (PGA).

You can control cache size for individual caches or globally, and automatically, by an Oracle database. Setting a unified memory size parameter enables Oracle to automatically manage individual cache sizes.

- All Oracle memory parameters are set using the `ALTER SYSTEM` command.
- Some changes to memory parameters require an instance restart.

Some of the common Oracle parameters that control memory allocations include:

- `db_cache_size` — The size of the cache used for database data.
- `log_buffer` — The cache used to store Oracle redo log buffers until they are written to disk.
- `shared_pool_size` — The cache used to store shared cursors, stored procedures, control structures, and other structures.
- `large_pool_size` — The caches used for parallel queries and RMAN backup/restore operations.
- `java_pool_size` — The caches used to store Java code and JVM context.

While these parameters can be configured individually, most database administrators choose to let Oracle automatically manage RAM. Database administrators configure the overall size of the SGA, and Oracle sizes individual caches based on workload characteristics.

- `sga_max_size` — Specifies the hard-limit maximum size of the SGA.
- `sga_target` — Sets the required soft-limit for the SGA and the individual caches within it.

Oracle also provides control over how much private memory is dedicated for each session. Database Administrators configure the total size of memory available for all connecting sessions,

and Oracle allocates individual dedicated chunks from the total amount of available memory for each session.

- `pga_aggregate_target` — A soft-limit controlling the total amount of memory available for all sessions combined.
- `pga_aggregate_limit` — A hard-limit for the total amount of memory available for all sessions combined (Oracle 12c only).

In addition, instead of manually configuring the SGA and PGA memory areas, you can also configure one overall memory limit for both the SGA and PGA and let Oracle automatically balance memory between the various memory pools. This behavior is configured by the `memory_target` and `memory_max_target` parameters.

For more information, see [Memory Architecture](#) and [Database Memory Allocation](#) in the *Oracle documentation*.

MySQL usage

Such as other databases, MySQL uses different memory buffers for different purposes. In MySQL, there are several storage engines that use different memory buffers. This section refers to InnoDB only.

MySQL provides control over how server RAM is allocated. Some of the most important MySQL memory parameters include:

Memory pool parameter	Description
<code>innodb_buffer_pool_size</code>	The memory area where InnoDB caches table and index data.
<code>optimizer_trace_max_mem_size</code>	Buffer for optimizer traces.
<code>binlog_cache_size</code>	The size of the cache holding changes to the binary log during a transaction.
<code>host_cache_size</code>	Buffer area to store data on connections.
<code>innodb_ft_cache_size</code>	Very similar to <code>innodb_buffer_pool_size</code> but only for data related to <code>FULL_TEXT</code> indexes.

Memory pool parameter	Description
<code>stored_program_cache</code>	Cached stored routines per connection.
<code>sort_buffer_size</code>	Size of sort buffers used to sort data during creation of an InnoDB index.
Total memory available for a MySQL cluster	Controlled by selecting the DB Instance Class during instance creation: <div data-bbox="829 569 1507 1283" style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p>Instance specifications Estimate your monthly costs for the DB Instance using the Simple Monthly Calculator.</p> <p>DB engine Aurora - compatible with MySQL 5.7.12</p> <p>DB instance class info</p> <div style="border: 1px solid #ccc; padding: 5px;"> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.t2.small — 1 vCPU, 2 GiB RAM ▾</div> <div style="border-bottom: 1px solid #ccc; padding: 2px; background-color: #e6f2ff;">db.t2.small — 1 vCPU, 2 GiB RAM</div> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.t2.medium — 2 vCPU, 4 GiB RAM</div> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.r3.large — 2 vCPU, 15.25 GiB RAM</div> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.r3.xlarge — 4 vCPU, 30.5 GiB RAM</div> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.r3.2xlarge — 8 vCPU, 61 GiB RAM</div> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.r3.4xlarge — 16 vCPU, 122 GiB RAM</div> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.r3.8xlarge — 32 vCPU, 244 GiB RAM</div> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.r4.large — 2 vCPU, 15.25 GiB RAM</div> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.r4.xlarge — 4 vCPU, 30.5 GiB RAM</div> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.r4.2xlarge — 8 vCPU, 61 GiB RAM</div> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.r4.4xlarge — 16 vCPU, 122 GiB RAM</div> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.r4.8xlarge — 32 vCPU, 244 GiB RAM</div> <div style="border-bottom: 1px solid #ccc; padding: 2px;">db.r4.16xlarge — 64 vCPU, 488 GiB RAM</div> </div> </div>

Note

You can configure cluster-level parameters such as `innodb_buffer_pool_size` and `binlog_cache_size` using parameter groups in the Amazon Relational Database Service console.

Examples

View the configured values for database parameters.

```
SHOW VARIABLES LIKE 'innodb_buffer_pool_size';
SHOW VARIABLES LIKE 'binlog_cache_size';
SHOW VARIABLES LIKE 'stored_program_cache';
```

View the configured values for all database parameters.

```
SELECT * FROM information_schema.GLOBAL_VARIABLES
```

Use the SET SESSION command to modify the value of parameters that support session-specific settings. Changing the value for one session has no effect on other sessions.

```
SET SESSION sort_buffer_size = 1000000;
```

Summary

Use the table below as a general reference only. Functionality may not be identical across Oracle and MySQL.

Description	Oracle	MySQL
Memory for caching table data	db_cache_size	innodb_buffer_pool_size
Memory for transaction log records	log_buffer	binlog_cache_size
Memory for parallel queries	large_pool_size	N/A
Java code and JVM	Java_pool_size	N/A
Maximum amount of physical memory available for the instance	sga_max_size or memory_max_size	Configured by the Amazon Relational Database Service or Amazon Aurora instance class. Consider the following example: db.r3.large: 15.25GB

Description	Oracle	MySQL
		db.r3.xlarge: 30.5GB
Total amount of private memory for all sessions	pga_aggregate_target and pga_aggregate_limit	max_digest_length
View values for all database parameters	<pre>SELECT * FROM v\$parameter;</pre>	<pre>SELECT * FROM information_schema.GLOBAL_VARIABLES</pre>
Configure a session-level parameter	<pre>ALTER SESSION SET ...</pre>	<pre>SET SESSION ...</pre>
Configure instance-level parameter	<pre>ALTER SYSTEM SET ...</pre>	Configured through parameter groups in the Amazon Relational Database Service console.

For more information, see [InnoDB Startup Options and System Variables](#) and [How MySQL Uses Memory](#) in the *MySQL documentation*.

Oracle instance parameters and Aurora MySQL parameter groups

With Amazon DMS, you can configure database settings to optimize performance, security, and resource utilization during and after migrating databases to Amazon Aurora. Oracle instance parameters and Aurora MySQL parameter groups define settings that govern database behavior, such as memory allocation, query optimization, and security policies.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Use cluster and database cluster parameters.

Oracle usage

You can configure Oracle instance and database-level parameters using the `ALTER SYSTEM` command. You can configure certain parameters dynamically and take immediate effect while other parameters require an instance restart.

- All Oracle instance and database-level parameters are stored in a binary file known as the Server Parameter file (SPFILE).
- The binary SPFILE can be exported to a text file using the following command:

```
CREATE PFILE = 'my_init.ora'
FROM SPFILE = 's_params.ora';
```

When you modify parameters, you can choose the persistence of the changed values with one of the three following options:

- Make the change applicable only after a restart by specifying `scope=spfile`.
- Make the change dynamically, but not persistent, after a restart by specifying `scope=memory`.
- Make the change both dynamically and persistent by specifying `scope=both`.

Examples

Use the `ALTER SYSTEM SET` command to configure a value for an Oracle parameter.

```
ALTER SYSTEM SET QUERY_REWRITE_ENABLED = TRUE SCOPE=BOTH;
```

For more information, see [Initialization Parameters](#) and [Changing Parameter Values in a Parameter File](#) in the *Oracle documentation*.

MySQL usage

When you run MySQL databases as Amazon Aurora clusters, you can use parameter groups to change the cluster-level and database-level parameters.

Most of the MySQL parameters are configurable in an Amazon Aurora MySQL cluster, but some are disabled and cannot be modified. Since Amazon Aurora clusters restrict access to the underlying operating system, modification to MySQL parameters must be made using Parameter Groups.

Amazon Aurora is a cluster of database instances and, as a direct result, some of the MySQL parameters apply to the entire cluster while other parameters apply only to a particular database instance.

Aurora MySQL parameter class	Controlled through
<p>Cluster-level parameters</p> <p>Single cluster parameter group for each Amazon Aurora cluster</p>	<p>Managed using cluster parameter groups.</p> <p>Consider the following example:</p> <pre data-bbox="829 1129 1507 1287">aurora_load_from_s3_role, default_password_lifetime, default_storage_engine</pre>
<p>Database instance-level parameters</p> <p>Every instance in an Amazon Aurora cluster can be associated with a unique database parameter group</p>	<p>Managed using database parameter groups.</p> <p>Consider the following example:</p> <pre data-bbox="829 1480 1507 1638">autocommit, connect_timeout, innodb_change_buffer_max_size</pre>

Examples

Create and configure a new parameter group

Follow these steps to create and configure an Amazon Aurora database and cluster parameter groups:

1. Sign in to the Amazon Management Console, choose **RDS**, and then choose **Databases**.
2. Choose **Parameter groups**, and choose **Create parameter group**.

 **Note**

You can't edit the default parameter group. Create a custom parameter group to apply changes to your Amazon Aurora cluster and its database instances.

3. For **Parameter group family**, choose **aurora-mysql5.7**.
4. For **Type**, choose **DB parameter group**.
5. Choose **Create**.

Modify an existing parameter group

1. Sign in to the Amazon Management Console, choose **RDS**, and then choose **Databases**.
2. Choose **Parameter groups**, and choose the name of the parameter to edit.
3. For **Parameter group actions**, choose **Edit**.
4. Change parameter values and choose **Save changes**.

For more information, see [Server System Variables](#) in the *MySQL documentation*.

Oracle session parameters and MySQL session variables

With Amazon DMS, you can configure Oracle session parameters and MySQL session variables to optimize performance, control resource usage, and customize database behavior during migration tasks. Oracle session parameters and MySQL session variables are special configuration settings that influence how the database engine operates and processes data. These settings can be crucial for ensuring efficient data transfer, minimizing resource contention, and adhering to organizational policies or regulatory requirements.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	SET options are significantly different.

Oracle usage

Certain Oracle database parameters and configuration options are modifiable at the session level using the ALTER SESSION command. However, not all Oracle configuration options and parameters can be modified on a per-session basis. To view a list of all configurable parameters that can be set for the scope of a specific session, query the v\$parameter view as shown in the following example.

```
SELECT NAME, VALUE FROM V$PARAMETER WHERE ISSES_MODIFIABLE='TRUE';
```

Examples

Change the NLS_LANGUAGE codepage parameter of the current session.

```
alter session set nls_language='SPANISH'

Sesión modificada.

alter session set nls_language='ENGLISH';

Session altered.

alter session set nls_language='FRENCH';

Session modifiée.

alter session set nls_language='GERMAN';

Session wurde geändert.
```

Specify the format of date values returned from the database using the `NLS_DATE_FORMAT` session parameter.

```
select sysdate from dual;

SYSDATE
SEP-09-17

alter session set nls_date_format='DD-MON-RR';
Session altered.

select sysdate from dual;

SYSDATE
09-SEP-17

alter session set nls_date_format='MM-DD-YYYY';
Session altered.

select sysdate from dual;

SYSDATE
09-09-2017

alter session set nls_date_format='DAY-MON-RR';
Session altered.
```

For more information, see [Changing Parameter Values in a Parameter File](#) in the *Oracle documentation*.

MySQL usage

MySQL provides session-modifiable parameters configured using the `SET SESSION` command. Configuration of parameters using `SET SESSION` is only applicable in the current session. To view the list of parameters that you can set with `SET SESSION`, see [Dynamic System Variables](#) and search for variables with session scope.

Examples of commonly used session parameters:

- `autocommit` — Specify if changes take effect immediately or if an explicit `COMMIT` command is required.

- `character_set_client` — Set the character set for the client.
- `default_storage_engine` — Set the default storage engine.
- `foreign_key_checks` — Set whether or not to run FK checks.
- `innodb_lock_wait_timeout` — Set how much time the transaction should wait to acquire a row lock.

Examples

Change the time zone of the connected session.

```
SELECT now();

now()
2018-02-26 12:13:25

SET SESSION TIME_ZONE = '+10:00';
SELECT now();

now()
2018-02-26 22:14:03
```

You can also use a time zone name such as `Europe/Helsinki` instead of `+10:00`.

Oracle and MySQL session parameter examples

Parameter purpose	Oracle	MySQL
Configure time and date format	<pre>ALTER SESSION SET nls_date_format = 'dd/mm/yyyy hh24:mi:ss';</pre>	N/A
Configure the current default schema or database	<pre>ALTER SESSION SET current schema='s chema_name'</pre>	N/A
Generate traces for specific errors	<pre>ALTER SESSION</pre>	N/A

Parameter purpose	Oracle	MySQL
	<pre>SET events '10053 trace name context forever';</pre>	
Run trace for a SQL statement	<pre>ALTER SESSION SET sql_trace=TRUE; ALTER SYSTEM SET EVENTS 'sql_trace [sql:&&sql_id] bind=true, wait=true';</pre>	<pre>SET GLOBAL general_log = 'ON';</pre>
Modify query optimizer cost for index access	<pre>ALTER SESSION SET optimizer_index_co st_adj = 50</pre>	<pre>SET SESSION optimizer _switch= ?</pre> <p>You can turn on and off other strategies. For more information, see Switchable Optimizations in the <i>MySQL documentation</i>.</p>
Modify query optimizer row access strategy	<pre>ALTER SESSION SET optimizer_mode=all _rows;</pre>	<pre>SET SESSION optimizer _switch= ?</pre> <p>You can turn on and off other strategies. For more information, see Switchable Optimizations in the <i>MySQL documentation</i>.</p>
Memory allocated to sort operations	<pre>ALTER SESSION SET sort_area_size=632 1;</pre>	<pre>SET SESSION sort_buff er_size=32768;</pre>

Parameter purpose	Oracle	MySQL
Memory allocated to hash-joins	<pre>ALTER SESSION SET hash_area_size= 1048576000;</pre>	<pre>SET SESSION join_buffer_size=1048576000;</pre>

For more information, see [SET Syntax for Variable Assignment](#) in the *MySQL documentation*.

Oracle and MySQL performance tuning

This section provides pages related to Oracle and MySQL performance tuning topics.

Topics

- [Database hints](#)
- [Oracle and MySQL run plans](#)
- [Oracle table statistics and MySQL managing statistics](#)

Database hints

With Amazon DMS, you can configure Oracle session parameters and MySQL session variables to optimize performance, control resource usage, and customize database behavior during migration tasks. Oracle session parameters and MySQL session variables are special configuration settings that influence how the database engine operates and processes data. These settings can be crucial for ensuring efficient data transfer, minimizing resource contention, and adhering to organizational policies or regulatory requirements.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Very limited set of hints in MySQL. Use index hints and optimizer hints as comments. Syntax differences.

Oracle usage

Oracle provides users with the ability to influence how the query optimizer behaves and the decisions made to generate query run plans. Controlling the behavior of the database optimizer is performed using database hints. They can be defined as a directive operation to the optimizer and alter the decisions of how run plans are generated.

Oracle supports over 60 different database hints, and each database hint can have 0 or more arguments. Database hints are divided into different categories such as optimizer hints, join order hints, and parallel run hints.

Database hints are embedded directly into the SQL queries immediately following the SELECT keyword using the following format: `/* <DB_HINT> */`.

Examples

Force the query optimizer to use a specific index for data access.

```
SELECT /* INDEX(EMP, IDX_EMP_HIRE_DATE)*/ * FROM EMPLOYEES EMP
WHERE HIRE_DATE >= '01-JAN-2010';
```

Run Plan

Plan hash value: 3035503638

Id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time
0	SELECT STATEMENT		1	62	2 (0)	
00:00:01						
1	TABLE ACCESS BY INDEX ROWID	EMPLOYEES	1	62	2 (0)	
00:00:01						
* 2	INDEX RANGE SCAN	IDX_HIRE_DATE	1		1 (0)	
00:00:01						

Predicate Information (identified by operation id):

2 - access("HIRE_DATE">=TO_DATE(' 2010-01-01 00:00:00', 'syyyy-mm-dd hh24:mi:ss'))

For more information, see [Additional Hints](#) and [Influencing the Optimizer](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL supports two types of hints: optimizer hints and index hints.

Index hints

The `USE INDEX` hint limits the optimizer's choice to one of the indexes listed in the `<Index List>` white list. Alternatively, indexes can be black listed using the `IGNORE` keyword.

The `FORCE INDEX` hint is similar to `USE INDEX (index_list)`, but with strong favor towards seek against scan. The hints use the actual index names, not column names. You can refer to primary keys using the keyword `PRIMARY`.

Syntax

```
SELECT ...
FROM <Table Name>
    USE {INDEX|KEY}
        [FOR {JOIN|ORDER BY|GROUP BY}] (<Index List>)
    | IGNORE {INDEX|KEY}
        [FOR {JOIN|ORDER BY|GROUP BY}] (<Index List>)
    | FORCE {INDEX|KEY}
        [FOR {JOIN|ORDER BY|GROUP BY}] (<Index List>)
...n
```

Note

In Aurora MySQL, the primary key is the clustered index.

The syntax for index hints has the following characteristics: * You can omit <Index List> for USE INDEX only. It translates to *don't use any indexes*, which is equivalent to a clustered index scan. * Index hints can be further scoped down using the FOR clause. Use FOR JOIN, FOR ORDER BY, or FOR GROUP BY to limit the hint applicability to that specific query processing phase. * Multiple index hints can be specified for the same or different scope.

Optimizer hints

Optimizer hints give developers or administrators control over some of the optimizer decision tree. They are specified within the statement text as a comment with the prefix +.

Optimizer hints may pertain to different scopes and are valid in only one or two scopes. The available scopes for optimizer hints in descending scope width order are:

- **Global** hints affect the entire statement. Only MAX_EXECUTION TIME is a global optimizer hint.
- **Query-level** hints affect a query block within a composed statement such as UNION or a subquery.
- **Table-level** hints affect a table within a query block.
- **Index-level** hints affect an index of a table.

Syntax

```
SELECT /*+ <Optimizer Hints> */ <Select List>...
```

```
INSERT /*+ <Optimizer Hints> */ INTO <Table>...
```

```
REPLACE /*+ <Optimizer Hints> */ INTO <Table>...
```

```
UPDATE /*+ <Optimizer Hints> */ <Table> SET...
```

```
DELETE /*+ <Optimizer Hints> */ FROM <Table>...
```

The following optimizer hints are available in Aurora MySQL.

Hint Name	Description	Applicable Scopes
BKA, NO_BKA	Turns on or turns off batched key access join processing	Query block, table
BNL, NO_BNL	Turns on or turns off block nested loop join processing	Query block, table
MAX_EXECUTION_TIME	Limits statement run time	Global
MRR, NO_MRR	Turns on or turns off multi-range read optimization	Table, index
NO_ICP	Turns off index condition push-down optimization	Table, index
NO_RANGE_OPTIMIZATION	Turns off range optimization	Table, index
QB_NAME	Assigns a logical name to a query block	Query block
SEMIJOIN, NO_SEMIJOIN	Turns on or turns off semi-join strategies	Query block

Hint Name	Description	Applicable Scopes
SUBQUERY	Determines MATERIALI ZATION , and INTOEXISTS processing	Query block

You can use query block names with QB_NAME to distinguish a block for limiting the scope of the table hint. Add @ to indicate a hint scope for one or more named subqueries. Consider the following example:

```
SELECT /*+ SEMIJOIN(@SubQuery1 FIRSTMATCH, LOOSESCAN) */ *
FROM Table1
WHERE Col1 IN (SELECT /*+ QB_NAME(SubQuery1) */ Col1
FROM t3);
```

Values for MAX_EXECUTION_TIME are measured in seconds and are always global for the entire query.

Note

This option doesn't exist in Oracle, where the run time limit pertains to the session scope.

Migration considerations

In general, the Aurora MySQL hint framework is relatively limited compared to the granular control provided by Oracle. It is recommended to start migration testing with all hints removed. Then, selectively apply hints as a last resort if other means such as schema, index, and query optimizations have failed.

Aurora MySQL uses a list of indexes and hints, both white list (USE) and black list (IGNORE), as opposed to Oracle's explicit index approach.

Index hints are not mandatory instructions. Aurora MySQL may choose alternatives if it cannot use the hinted index.

Examples

Force an index access.

```
SELECT * FROM Table1 USE INDEX (Index1) ORDER BY Col1;
```

Specify multiple index hints.

```
SELECT * FROM Table1
  USE INDEX (Index1)
 INNER JOIN Table2
   IGNORE INDEX(Index2)
   ON Table1.Col1 = Table2.Col1
 ORDER BY Col1;
```

Specify optimizer hints.

```
SELECT /*+ NO_RANGE_OPTIMIZATION(Table1 PRIMARY, Index2) */
Col1 FROM Table1 WHERE Col2 = 300;
```

```
SELECT /*+ BKA(t1) NO_BKA(t2) */ * FROM Table1 INNER JOIN Table2 ON ...;
```

```
SELECT /*+ NO_ICP(t1, t2) */ * FROM Table1 INNER JOIN Table2 ON ...;
```

Summary

Feature	Oracle	Aurora MySQL
Force a specific plan	DBMS_SPM	N/A
Join hints	USE_NL, NO_USE_NL , USE_NL_WITH_INDEX , USE_MERGE , NO_USE_MERGE , USE_HASH, NO_USE_HASH	BNL, NO_BNL (Block Nested Loops)
Force scan	FULL	USE with no index list forces a clustered index scan
Force an index	INDEX	USE

Feature	Oracle	Aurora MySQL
Allow list and deny list indexes	NO_INDEX	Supported with USE and IGNORE
Parameter value hints	opt_param	N/A

For more information, see [Controlling the Query Optimizer](#), [Optimizer Hints](#), [Index Hints](#), and [Optimizing Subqueries, Derived Tables, and View References](#) in the *MySQL documentation*.

Oracle and MySQL run plans

With Amazon DMS, you can analyze the performance of your database migration tasks by reviewing Oracle and MySQL run plans. An Oracle or MySQL run plan provides detailed information about the run plan for a specific SQL statement, including the steps involved, data access methods, and potential performance bottlenecks.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Syntax differences. Completely different optimizer with different operators and rules.

Oracle usage

Run plans represent the choices made by the query optimizer for accessing database data. The query optimizer generates run plans for SELECT, INSERT, UPDATE, and DELETE statements. Users and database administrators can view run plans for specific queries and DML operations.

Run plans are especially useful for performance tuning of queries. For example, determining if new indexes should be created. Run plans can be affected by data volumes, data statistics, and instance parameters such as global or session-level parameters.

Run plans are displayed as a structured tree with the following information:

- Tables access by the SQL statement and the referenced order for each table.
- Access method for each table in the statement such as full table scan or index access.
- Algorithms used for join operations between tables such as hash or nested loop joins.
- Operations performed on retrieved data as such as filtering, sorting, and aggregations.
- Information about rows being processed (cardinality) and the cost for each operation.
- Table partitions being accessed.
- Information about parallel runs.

Oracle 19 introduces SQL Quarantine: now queries that consume resources excessively can be automatically quarantined and prevented from running. These queries run plans are also quarantined.

Examples

Review the potential run plan for a query using the EXPLAIN PLAN statement.

```

SET AUTOTRACE TRACEONLY EXPLAIN
SELECT EMPLOYEE_ID, LAST_NAME, FIRST_NAME FROM EMPLOYEES
WHERE LAST_NAME='King' AND FIRST_NAME='Steven';
Run Plan

Plan hash value: 2077747057
| Id | Operation                | Name          | Rows | Bytes | Cost (%CPU) | Time
| 0  | SELECT STATEMENT         |               | 1    | 16    | 2 (0)       |
00:00:01
| 1  | TABLE ACCESS BY INDEX ROWID | EMPLOYEES    | 1    | 16    | 2 (0)       |
00:00:01
|* 2  | INDEX RANGE SCAN         | EMP_NAME_IX  | 1    |       | 1 (0)       |
00:00:01

Predicate Information (identified by operation id):
2 - access("LAST_NAME"='King' AND "FIRST_NAME"='Steven')

```

SET AUTOTRACE TRACEONLY EXPLAIN instructs SQL*PLUS to show the run plan without actually running the query itself.

The EMPLOYEES table contains indexes for both the LAST_NAME and FIRST_NAME columns. Step 2 of the run plan above indicates the optimizer is performing an INDEX RANGE SCAN to retrieve the filtered employee name.

View a different run plan displaying a FULL TABLE SCAN.

```
SET AUTOTRACE TRACEONLY EXPLAIN
SELECT EMPLOYEE_ID, LAST_NAME, FIRST_NAME FROM EMPLOYEES
WHERE SALARY > 10000;
Run Plan

Plan hash value: 1445457117
| Id | Operation          | Name          | Rows | Bytes | Cost (%CPU) | Time
| 0  | SELECT STATEMENT   |               | 72   | 1368 | 3 (0)        | 00:00:01
|* 1 | TABLE ACCESS FULL| EMPLOYEES    | 72   | 1368 | 3 (0)        | 00:00:01

Predicate Information (identified by operation id):
1 - filter("SALARY">10000)
```

For more information, see [Explaining and Displaying Execution Plans](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL provides the EXPLAIN/DESCRIBE statement—used with the SELECT, DELETE, INSERT, REPLACE, and UPDATE statements—to display run plans.

Note

You can use the EXPLAIN/DESCRIBE statement to retrieve table and column metadata.

When you use EXPLAIN with a statement, MySQL returns the run plan generated by the query optimizer. MySQL explains how the statement will be processed including information about table joins and order.

When you use EXPLAIN with the FOR CONNECTION option, it returns the run plan for the statement running in the named connection. You can use the FORMAT option to select either a TRADITIONAL tabular format or JSON.

The EXPLAIN statement requires SELECT permissions for all tables and views accessed by the query directly or indirectly. For views, EXPLAIN requires the SHOW VIEW permission.

EXPLAIN can be extremely valuable for improving query performance when used to find missing indexes. You can also use EXPLAIN to determine if the optimizer joins tables in an optimal order.

MySQL Workbench includes an easy to read visual explain feature similar to Oracle Execution Manager (OEM) graphical run plans.

Note

Amazon Relational Database Service (Amazon RDS) for MySQL version 8.0.18 implements EXPLAIN ANALYZE, a new form of the EXPLAIN statement. This statement provides expanded information about the run of SELECT statements in the TREE format for each iterator used in processing the query and making it possible to compare estimated cost with the actual cost of the query. This information includes startup cost, total cost, number of rows returned by this iterator and the number of loops executed. In MySQL 8.0.21 and later, this statement also supports a FORMAT=TREE specifier. TREE is the only supported format. For more information, see [Obtaining Information with EXPLAIN ANALYZE](#) in the *MySQL documentation*.

Syntax

The following example shows the simplified syntax for the EXPLAIN statement.

```
{EXPLAIN | DESCRIBE | DESC} [EXTENDED | FORMAT = TRADITIONAL | JSON]
[SELECT statement | DELETE statement | INSERT statement | REPLACE statement | UPDATE
statement | FOR CONNECTION <connection id>]
```

Examples

View the run plan for a statement.

```
CREATE TABLE Employees (
  EmployeeID INT NOT NULL PRIMARY KEY,
  Name VARCHAR(100) NOT NULL,
  INDEX USING BTREE(Name));
```

```
EXPLAIN SELECT * FROM Employees WHERE Name = 'Jason';
```

For the preceding example, the result looks as shown following.

id	select_type	table	partitions	type	possible_keys	key	key_len	ref	rows
Extra									
1	SIMPLE	Employees		ref	Name	Name	102		const
1									

The following image demonstrates the MySQL Workbench graphical run plan.

The screenshot displays the MySQL Workbench interface. At the top, the query editor shows the following SQL statement:

```
1 SELECT * FROM orders
2 WHERE o_orderdate BETWEEN '1992-04-01' AND '1992-04-30'
3 AND o_clerk LIKE '%0223';
```

Below the query editor, the 'Visual Explain' tab is active, showing a graphical run plan. The plan consists of a single node: 'query_block #1', which leads to an 'Index Range Scan' operation on the 'orders' table, specifically using the 'i_o_orderdate' index. A tooltip for this operation provides the following details:

- Access Type: range
- Index Range Scan
- Cost Hint: Medium - partial index scan
- Key/Index: i_o_orderdate
- Used Key Parts: o_orderDATE
- Possible Keys: i_o_orderdate
- Attached Condition: ('dbt3`.`orders`.`o_clerk` like '%0223')
- Rows Examined per Scan: 32642
- Rows Produced per Join: 32642
- Filtered (ratio of rows produced per rows examined): 100%
- Hint: 100% is best, <= 1% is worst
- A low value means the query examines a lot of rows that are not returned.

At the bottom of the window, the 'Output' tab shows the execution results:

Time	Action	Status	Details
00:00:44	SELECT * FROM orders WHERE o_orderdate BETWEEN '1992-04-01' AND '1992-04-30' AND o_clerk LIKE '%0223' LIMIT 0, 1000	OK	18 row(s) return... 0.281 sec / 0.000 sec
00:00:49	EXPLAIN SELECT * FROM orders WHERE o_orderdate BETWEEN '1992-04-01' AND '1992-04-30' AND o_clerk LIKE '%0223'	OK	0.000 sec

Note

To instruct the optimizer to use a join order corresponding to the order in which the tables are specified in a SELECT statement, use `SELECT STRAIGHT_JOIN`.

For more information, see [EXPLAIN Statement](#) in the *MySQL documentation*.

Oracle table statistics and MySQL managing statistics

With Amazon DMS, you can gather and manage statistics about database tables and indexes to improve query performance. Oracle table statistics and MySQL managing statistics provide mechanisms to collect and update metadata about the distribution of data in tables and associated indexes. This information helps the query optimizer generate efficient run plans.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Syntax and option differences, similar functionality.

Oracle usage

Table statistics are one of the important aspects affecting SQL query performance. They turn on the query optimizer to make informed assumptions when deciding how to generate the run plan for each query. Oracle provides the DBMS_STATS package to manage and control the table statistics, which you can collect automatically or manually.

The following statistics are usually collected on database tables and indexes:

- Number of table rows.
- Number of table blocks.
- Number of distinct values or nulls.
- Data distribution histograms.

Automatic optimizer statistics collection

By default, Oracle collects table and index statistics during predefined maintenance windows using the database scheduler and automated maintenance tasks. The automatic statistics collection mechanism uses Oracle data modification monitoring feature that tracks the approximate number of INSERT, UPDATE, and DELETE statements to determine which table statistics should be collected.

In Oracle 19, you can gather real-time statistics on tables during regular UPDATE, INSERT, and DELETE operations, which ensures that statistics are always up-to-date and are not going stale.

Oracle 19 also introduces high-frequency automatic optimizer statistics collection. Use this feature to set up automatic task that will collect statistics for stale objects.

Manual optimizer statistics collection

When the automatic statistics collection is not suitable for a particular use case, you can perform the optimizer statistics collection manually at several levels:

Statistics level	Description
GATHER_INDEX_STATS	Index statistics
GATHER_TABLE_STATS	Table, column, and index statistics
GATHER_SCHEMA_STATS	Statistics for all objects in a schema
GATHER_DICTIONARY_STATS	Statistics for all dictionary objects
GATHER_DATABASE_STATS	Statistics for all objects in a database

Examples

Collect statistics at the table level for the HR schema and the EMPLOYEES table.

```
BEGIN
DBMS_STATS.GATHER_TABLE_STATS('HR', 'EMPLOYEES');
END;
/
```

PL/SQL procedure successfully completed.

Collect statistics at a specific column level for the HR schema, the EMPLOYEES table, and the DEPARTMENT_ID column.

```
BEGIN
DBMS_STATS.GATHER_TABLE_STATS('HR', 'EMPLOYEES',
```

```
METHOD_OPT=>'FOR COLUMNS department_id');
END;
/
```

PL/SQL procedure successfully completed.

For more information, see [Optimizer Statistics Concepts](#) in the *Oracle documentation*.

MySQL usage

Aurora MySQL supports two modes of statistics management: Persistent Optimizer Statistics and Non-Persistent Optimizer Statistics. As the name suggests, persistent statistics are written to disk and survive service restart. Non-persistent statistics are kept in memory and need to be recreated after service restart. It is recommended to use persistent optimizer statistics (the default for Aurora MySQL) for improved plan stability.

Statistics in Aurora MySQL are created for indexes only. Aurora MySQL does not support independent statistics objects on columns that are not part of an index.

Typically, administrators change the statistics management mode by setting the global parameter `innodb_stats_persistent = ON`. Therefore, control the statistics management mode by changing the behavior for individual tables using the table option `STATS_PERSISTENT = 1`. There are no column-level or statistics-level options for setting parameter values.

To view statistics metadata, use the `INFORMATION_SCHEMA.STATISTICS` standard view. To view detailed persistent optimizer statistics, use the `innodb_table_stats` and `innodb_index_stats` tables.

The following image demonstrates an example of the `mysql.innodb_table_stats` content.

database_name	table_name	last_update	n_rows	clustered_index_size	sum_of_other_index_sizes
mysql	aurora_s3_load_history	2021-06-16 01:44:29	0	1	0
mysql	columns_priv	2017-07-14 00:25:33	0	1	0
mysql	db	2017-07-14 00:25:33	0	1	1
mysql	event	2017-10-31 02:20:35	0	1	0
mysql	func	2017-07-14 00:25:33	0	1	0
mysql	host	2017-10-26 21:31:42	0	1	0
mysql	ndb_binlog_index	2017-07-13 18:08:32	0	1	0
mysql	proc	2021-06-16 01:44:40	79	97	0
mysql	procs_priv	2017-10-31 02:20:25	0	1	1
mysql	proxies_priv	2017-07-13 18:08:32	0	1	1
mysql	rds_configuration	2017-07-13 18:08:32	0	1	0
mysql	rds_global_status_history	2017-07-13 18:08:32	0	1	0
mysql	rds_global_status_history_old	2017-07-13 18:08:32	0	1	0
mysql	rds_heartbeat2	2017-07-13 18:08:32	0	1	0
mysql	rds_history	2021-06-16 01:44:30	1	1	0
mysql	rds_replication_status	2017-10-31 02:18:52	0	1	0
mysql	rds_sysinfo	2017-07-13 18:08:32	0	1	0

The following image demonstrates an example of the `mysql.innodb_index_stats` content.

database_name	table_name	index_name	last_update	stat_name	stat_value	sample_size	stat_description
mysql	aurora_s3_load_history	GEN_CLUST_INDEX	2021-06-16 01:44:29	n_diff_pfx01	0	1	DB_ROW_ID
mysql	aurora_s3_load_history	GEN_CLUST_INDEX	2021-06-16 01:44:29	n_leaf_pages	1	1	Number of leaf pages in the index
mysql	aurora_s3_load_history	GEN_CLUST_INDEX	2021-06-16 01:44:29	size	1	1	Number of pages in the index
mysql	columns_priv	PRIMARY	2017-07-14 00:25:33	n_diff_pfx01	0	1	Host
mysql	columns_priv	PRIMARY	2017-07-14 00:25:33	n_diff_pfx02	0	1	Host,Db
mysql	columns_priv	PRIMARY	2017-07-14 00:25:33	n_diff_pfx03	0	1	Host,Db,User
mysql	columns_priv	PRIMARY	2017-07-14 00:25:33	n_diff_pfx04	0	1	Host,Db,User,Table_name
mysql	columns_priv	PRIMARY	2017-07-14 00:25:33	n_diff_pfx05	0	1	Host,Db,User,Table_name,Column_name
mysql	columns_priv	PRIMARY	2017-07-14 00:25:33	n_leaf_pages	1	1	Number of leaf pages in the index
mysql	columns_priv	PRIMARY	2017-07-14 00:25:33	size	1	1	Number of pages in the index
mysql	db	PRIMARY	2017-07-14 00:25:33	n_diff_pfx01	0	1	Host
mysql	db	PRIMARY	2017-07-14 00:25:33	n_diff_pfx02	0	1	Host,Db
mysql	db	PRIMARY	2017-07-14 00:25:33	n_diff_pfx03	0	1	Host,Db,User
mysql	db	PRIMARY	2017-07-14 00:25:33	n_leaf_pages	1	1	Number of leaf pages in the index
mysql	db	PRIMARY	2017-07-14 00:25:33	size	1	1	Number of pages in the index
mysql	db	User	2017-07-14 00:25:33	n_diff_pfx01	0	1	User
mysql	db	User	2017-07-14 00:25:33	n_diff_pfx02	0	1	User,Host
mysql	db	User	2017-07-14 00:25:33	n_diff_pfx03	0	1	User,Host,Db
mysql	db	User	2017-07-14 00:25:33	n_leaf_pages	1	1	Number of leaf pages in the index
mysql	db	User	2017-07-14 00:25:33	size	1	1	Number of pages in the index
mysql	event	PRIMARY	2017-10-31 02:20:35	n_diff_pfx01	0	1	db

Automatic refresh of statistics is controlled by the global parameter `innodb_stats_auto_recalc`, which is set to `ON` in Aurora MySQL. You can set it individually for each table using the `STATS_AUTO_RECALC=1` option.

To explicitly force a refresh of table statistics, use the `ANALYZE TABLE` statement. It is not possible to refresh individual statistics or columns.

Use the `NO_WRITE_TO_BINLOG`, or its clearer alias `LOCAL`, to avoid replication to replication secondaries.

Use `ALTER TABLE ... ANALYZE PARTITION` to analyze one or more individual partitions.

Note

Amazon Relational Database Service (Amazon RDS) for MySQL version 8 adds new `INFORMATION_SCHEMA.INNODB_CACHED_INDEXES` table which reports the number of index pages cached in the InnoDB buffer pool for each index.

Syntax

```
ANALYZE [NO_WRITE_TO_BINLOG | LOCAL] TABLE <Table Name> [,...];
```

```
CREATE TABLE ( <Table Definition> ) | ALTER TABLE <Table Name>
STATS_PERSISTENT = <1|0>,
STATS_AUTO_RECALC = <1|0>,
```

```
STATS_SAMPLE_PAGES = <Statistics Sampling Size>;
```

Migration considerations

Unlike Oracle, Aurora MySQL collects only density information. It does not collect detailed key distribution histograms. This difference is critical for understanding execution plans and troubleshooting performance issues that are not affected by individual values used by query parameters.

Statistics collection is managed at the table level. You cannot manage individual statistics objects or individual columns. In most cases, that should not pose a challenge for successful migration.

Examples

The following example creates a table with explicitly set statistics options.

```
CREATE TABLE MyTable
(Co11 INT NOT NULL AUTO_INCREMENT,
Co12 VARCHAR(255),
DateCo1 DATETIME,
PRIMARY KEY (Co11),
INDEX IDX_DATE (DateCo1)
) ENGINE=InnoDB,
STATS_PERSISTENT=1,
STATS_AUTO_RECALC=1,
STATS_SAMPLE_PAGES=25;
```

The following example refreshes all statistics for MyTable1 and MyTable2.

```
ANALYZE TABLE MyTable1, MyTable2;
```

The following example changes the MyTable settings to use non-persistent statistics.

```
ALTER TABLE MyTable STATS_PERSISTENT=0;
```

Summary

The following table identifies Aurora MySQL features. All of the features are accessed in Oracle using the DBMS_STATS package.

Feature	Aurora MySQL	Comments
Column statistics	N/A	
Index statistics	Implicit with every index	Statistics are maintained automatically for every table index.
Refresh or update statistics	ANALYZE TABLE	Minimal scope in Aurora MySQL is the entire table. No control over individual statistics.
Auto create statistics	N/A	
Auto update statistics	Use the STATS_AUTO_RECALC table option	
Statistics sampling	Use the STATS_SAMPLE_PAGES table option	Can only use page number, not percentage for STATS_SAMPLE_PAGES .
Full scan refresh	N/A	Using a very large STATS_SAMPLE_PAGES may serve the same purpose.
Non-persistent statistics	Use the STATS_PERSISTENT=0 table option	

For more information, see [The INFORMATION_SCHEMA COLUMN_STATISTICS Table](#), [Configuring Persistent Optimizer Statistics Parameters](#), [Configuring Non-Persistent Optimizer Statistics Parameters](#), and [Configuring Optimizer Statistics for InnoDB](#) in the *MySQL documentation*.

Oracle and MySQL security

This section includes pages about Oracle and MySQL security-related topics.

Topics

- [Encrypted connections](#)
- [Oracle transparent data encryption and Amazon Aurora MySQL encryption and column encryption](#)
- [Oracle roles and MySQL privileges](#)
- [Oracle database users and MySQL users](#)

Encrypted connections

With Amazon DMS, you can secure data transmission between the replication instance and the source or target database by using encrypted connections. Encrypted connections provide a private encrypted tunnel for data transfer, protecting sensitive information from unauthorized access or interception.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	N/A

Oracle usage

Oracle Database supports encrypting incoming data out of the box using native Oracle Net Services. You can encode data that is sent to and from the server using Advanced Encryption Standard (AES) algorithm, ARIA(Academia, Research Institute, and Agency) algorithm, GOsudarstvennyy STandard (GOST) algorithm, Korea Information Security Agency SEED algorithm and Triple-DES encryption (3DES).

Algorithms can be specified in the `sqlnet.ora` file for the clients and servers.

For more information, see [Configuring Oracle Database Network Encryption and Data Integrity](#) in the *Oracle documentation*.

SSL/TLS connections to the Oracle database are supported starting with Oracle 12c in the standard edition.

For more information, see [SSL Connection to Oracle DB using JDBC, TLSv1.2, JKS or Oracle Wallets \(12.2 and lower\)](#) in the *Oracle Developers Blog*.

MySQL usage

MySQL supports encrypted connections between clients and the server using the TLS (Transport Layer Security) protocol. TLS is sometimes referred to as SSL (Secure Sockets Layer) but MySQL does not actually use the SSL protocol for encrypted connections because its encryption is weak.

OpenSSL 1.1.1 supports the TLS v1.3 protocol for encrypted connections.

Note

Amazon Relational Database Service (Amazon RDS) for MySQL version 8.0.16 and higher supports TLS v1.3 as well if both the server and client are compiled using OpenSSL 1.1.1 or higher. For more information, see [Encrypted Connection TLS Protocols and Ciphers](#) in the *MySQL documentation*.

Oracle transparent data encryption and Amazon Aurora MySQL encryption and column encryption

With Amazon DMS, you can migrate databases that use Oracle transparent data encryption or Amazon Aurora MySQL encryption and column encryption to maintain data security during and after the migration process. Oracle transparent data encryption and Aurora MySQL encryption and column encryption provide data-at-rest encryption to protect sensitive information stored in databases.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	For more information, see Encrypting Amazon RDS resources .

Oracle usage

Oracle uses Transparent Data Encryption (TDE) to encrypt data stored on media in order to provide data at rest protection. Although Oracle uses authentication, authorization, and auditing to secure data in the database, TDE provides additional security at the operating system level.

As the name implies, encryption operations are performed automatically and are transparent to client applications. However, TDE does not address data in transit, which must be handled by network security protocols.

Characteristics of TDE include:

- The `ADMINISTER KEY MANAGEMENT` system privilege is required to configure TDE.
- Data can be encrypted at the column level or the tablespace level.
- Key encryption is managed in the external TDE Master Encryption Module.
- There is one root key for each database.

Examples

Configure the root encryption key

Specify the location of the encryption wallet using the `ENCRYPTION_WALLET_LOCATION` parameter. Use one of the following options:

- Regular filesystem.
- Multiple databases share the same file.
- ASM file system.

- ASM disk group.

Register the key file in the ASM disk group.

```
ENCRYPTION_WALLET_LOCATION=  
  (SOURCE=  
    (METHOD=FILE)  
    (METHOD_DATA=  
      (DIRECTORY=+ASM_file_path_of_the_diskgroup)))
```

Create a software keystore

Use one of the following three types of software keystores:

- Password-based.
- Auto-login.
- Local auto-login.

Create a password-based software keystore. The user must have the ADMINISTER KEY MANAGEMENT or SYSKM privilege.

```
sqlplus c##sec_admin as syskm  
Enter password: password  
Connected.  
  
ADMINISTER KEY MANAGEMENT CREATE KEYSTORE '/etc/ORACLE/WALLETS/orcl' IDENTIFIED BY  
password;  
  
keystore altered.
```

Open a keystore

When you use a password-based keystore, make sure that you open it before creating TDE master encryption keys or accessing the keystore. Keystores are automatically opened when using auto-login or local auto login.

```
sqlplus c##sec_admin as syskm  
Enter password: password  
Connected.
```

```
ADMINISTER KEY MANAGEMENT SET KEYSTORE OPEN IDENTIFIED BY password;  
  
keystore altered.
```

Set the software root encryption key

The master encryption key protects the TDE table and tablespace encryption keys. By default, the master encryption key is generated by TDE. To set the master encryption key, ensure the database is open in READ WRITE mode, connect with a user account having the required privileges (see the preceding example), and create the master key.

```
sqlplus c##sec_admin as syskm  
Enter password: password  
Connected.  
  
ADMINISTER KEY MANAGEMENT SET KEY IDENTIFIED BY keystore_password WITH BACKUP USING  
'emp_key_backup';  
  
keystore altered.
```

Encrypt data

Create an encrypted column.

```
CREATE TABLE employee (  
    FIRST_NAME VARCHAR2(128),  
    LAST_NAME VARCHAR2(128),  
    EMP_ID NUMBER,  
    SALARY NUMBER(6) ENCRYPT);
```

Column data types support for encryption include BINARY_DOUBLE, BINARY_FLOAT, CHAR, DATE, INTERVAL DAY TO SECOND, INTERVAL YEAR TO MONTH, NCHAR, NUMBER, NVARCHAR2, RAW (legacy or extended), TIMESTAMP (includes TIMESTAMP WITH TIME ZONE and TIMESTAMP WITH LOCAL TIME ZONE), VARCHAR2 (legacy or extended).

Column encryption can't be used with the following features:

- Index types other than B-tree.
- Range scan search through an index.
- Synchronous change data capture.

- Transportable tablespaces.
- Columns used in foreign key constraints.

You can change the encryption algorithm using the `NO SALT` clause to encrypt without an algorithm or the `USING` clause to specify an algorithm.

```
CREATE TABLE EMPLOYEE (  
    FIRST_NAME VARCHAR2(128),  
    LAST_NAME VARCHAR2(128),  
    EMP_ID NUMBER ENCRYPT NO SALT,  
    SALARY NUMBER(6) ENCRYPT USING '3DES168');
```

Change the algorithm on an existing table.

```
ALTER TABLE EMPLOYEE REKEY USING 'SHA-1';
```

Remove column encryption.

```
ALTER TABLE employee MODIFY (SALARY DECRYPT);
```

- Make sure that the `COMPATIBLE` initialization parameter is set to at least 11.2.0.0.
- Log in to your database.
- Create the tablespace. You can't modify an existing tablespace; you can only create a new one. In the following example, the first tablespace is created with AES256 algorithm and the second is created with the default algorithm.

```
sqlplus sec_admin@hrpdb  
Enter password: password  
Connected.  
  
CREATE TABLESPACE encrypt_ts  
DATAFILE '$ORACLE_HOME/dbs/encrypt_df.dbf' SIZE 1M  
ENCRYPTION USING 'AES256'  
DEFAULT STORAGE (ENCRYPT);  
CREATE TABLESPACE securespace_2  
DATAFILE '/home/user/oradata/secure01.dbf'  
SIZE 150M
```

```
ENCRYPTION  
DEFAULT STORAGE(ENCRYPT);
```

For more information, see [Introduction to Transparent Data Encryption](#) in the *Oracle documentation*.

MySQL usage

Amazon provides the ability to encrypt data at rest (data stored in persistent storage). When data encryption is turned on, it automatically encrypts the database server storage, automated backups, read replicas, and snapshots using the AES-256 encryption algorithm. Amazon Key Management Service (Amazon KMS) performs the encryption. For more information, see [Amazon Key Management Service](#).

Once enabled, Amazon transparently encrypts and decrypts the data without any impact on performance or any user intervention. There is no need to modify clients to support encryption.

Note

Amazon Relational Database Service (Amazon RDS) for MySQL version 8 supports FIPS mode if compiled using OpenSSL and an OpenSSL library and FIPS Object Module are available at runtime. FIPS mode imposes conditions on cryptographic operations such as restrictions on acceptable encryption algorithms or requirements for longer key lengths. For more information, see [FIPS Support](#) in the *MySQL documentation*.

Table encryption can now be managed globally by defining and enforcing encryption defaults. The `default_table_encryption` variable defines an encryption default for newly created schemas and general tablespaces. The encryption default for a schema can also be defined using the `DEFAULT ENCRYPTION` clause when creating a schema. By default a table inherits the encryption of the schema or general tablespace it is created in.

Encryption defaults are enforced by enabling the `table_encryption_privilege_check` variable. The privilege check occurs when creating or altering a schema or general tablespace with an encryption setting that differs from the `default_table_encryption` setting or when creating or altering a table with an encryption setting that differs from the default schema encryption. The `TABLE_ENCRYPTION_ADMIN` privilege permits overriding default encryption settings when `table_encryption_privilege_check` is enabled. For more information, see [Defining an Encryption Default for Schemas and General Tablespaces](#) in the *MySQL documentation*.

Create an encryption key

To create your own key, follow these steps.


1. Sign in to the Amazon Management Console and choose **Key Management Service**.
2. Choose **Customer managed keys**, and then choose **Create key**.
3. For **Key type**, choose **Symmetric**. Expand **Advanced options**. For **Key material origin**, choose **KMS**, and then choose **Next**.
4. For **Alias**, enter the name of your key. Choose **Next**.
5. On the **Define key administrative permissions** tab, choose **Next**.
6. On the next step, make sure that you assign the key to the relevant users who will need to interact with Amazon Aurora. Choose **Next**.
7. Review the key settings and choose **Finish** to create the key.
8. Set the Master encryption key. Use the ARN of the key that you created or choose this key from the list.

Now you can launch your instance.

Enabling encryption

As part of the database settings, you will be prompted to enable encryption and select a master key.

You can turn on encryption for an Amazon RDS DB instance only during the instance creation.



The screenshot shows the 'Encryption' configuration panel in the Amazon RDS console. Under the 'Encryption' heading, the 'Enable Encryption' radio button is selected. Below this, there is a 'Master key info' section with a dropdown menu currently showing '(default) aws/rds'. To the right of the main panel, a 'Master key' tooltip is displayed, stating: 'This is the master key that will be used to protect the key used to encrypt this database volume. You can select from master keys in your account or type/paste the ARN of a key from a different account. You can create a new master encryption key by going to the Encryption Keys tab of the IAM console.'

You can select the default key provided for the account or define a specific key based on an IAM KMS ARN from your account or a different account.

SSE-S3 encryption feature overview

Server-side encryption with Amazon S3-managed encryption keys (SSE-S3) uses multi-factor encryption. Amazon S3 encrypts its objects with a unique key and it also encrypts the key itself with a master key that rotates periodically.

SSE-S3 uses AES-256 as its encryption standard.

After you turn on the server-side encryption for an Amazon S3 bucket, the data will be encrypted at rest. Make sure that all API calls now include the special header as shown following: `-x-amz-server-side-encryption`.

For more information, see [Specifying Amazon S3 encryption](#) and [s3](#).

To turn on SSE-S3

1. Create an Amazon Glue job.
2. Define the role, bucket, and script and then open **Script libraries and job parameters (optional)**.
3. Turn on **Server-side encryption**.
4. Submit and run the job.

From this point forward, the only way to access the files is to use Amazon CLI `s3` with the `--sse` switch, or by adding `x-amz-server-side-encryption` to your API calls.

Usage of column encryption

Aurora MySQL provides encryption and decryption functions similar to Oracle with a much less elaborate security hierarchy that is easier to manage.

The encryption functions require the actual key as a string, so you must take extra measures to protect the data. For example, hashing the key values on the client.

Aurora MySQL supports the AES and DES encryption algorithms. You can use the following functions for data encryption and decryption:

- AES_DECRYPT
- AES_ENCRYPT
- DES_DECRYPT

- DEC_ENCRYPT

Syntax

General syntax for the encryption functions is shown following:

```
[A|D]ES_ENCRYPT(<string to be encrypted>, <key string> [,<initialization vector>])  
[A|D]ES_DECRYPT(<encrypted string>, <key string> [,<initialization vector>])
```

For more information, see [AES_ENCRYPT](#) in the *MySQL documentation*.

It is highly recommended to use the optional initialization vector to circumvent whole value replacement attacks. When encrypting column data, it is common to use an immutable key as the initialization vector. With this approach, decryption fails if a whole value moves to another row.

Consider using SHA2 instead of SHA1 or MD5 because there are known exploits available for the SHA1 and MD5. Passwords, keys, or any sensitive data passed to these functions from the client are not encrypted unless you are using an SSL connection. One benefit of using Amazon Identity and Access Management (IAM) is that database connections are encrypted with SSL by default. For more information, see [Users](#) and [Roles](#).

Examples

The following example demonstrates how to encrypt an employee social security number.

Create an employees table.

```
CREATE TABLE Employees (  
    EmployeeID INT NOT NULL PRIMARY KEY,  
    SSN_Encrypted BINARY(32) NOT NULL);
```

Insert the encrypted data.

```
INSERT INTO Employees (EmployeeID, SSN_Encrypted)  
VALUES (1, AES_ENCRYPT('1112223333', UNHEX(SHA2('MyPassword',512))), 1));
```

Note

Use the UNHEX function for more efficient storage and comparisons.

Verify decryption.

```
SELECT EmployeeID, SSN_Encrypted,
       AES_DECRYPT(SSN_Encrypted, UNHEX(SHA2('MyPassword', 512)), EmployeeID) AS SSN
FROM Employees
```

```
EmployeeID  SSN_Encrypted      SSN
1           ` ©> +yp°øýNZ~Gø  1112223333
```

For more information, see [Encryption and Compression Functions](#) in the *MySQL documentation*.

Oracle roles and MySQL privileges

With Amazon DMS, you can manage access control and security for your databases during migration. Oracle roles and MySQL privileges define permissions and access levels for database users, allowing you to restrict or grant specific operations and data access.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	There are no roles in MySQL, only privileges.

Oracle usage

Oracle roles are groups of privileges granted to database users. A database role can contain individual system and object permissions as well as other roles. Database roles enable you to grant multiple database privileges to users in one operation. It is convenient to group permissions together to ease the management of privileges.

Oracle 12c introduces a new multi-tenant database architecture that supports the creation of common and local roles:

- **Common** — Roles created at the container database (CDB) level. A common role is a database role that exists in the root and in every existing and future pluggable database (PDB). Common

roles are useful for cross-container operations such as ensuring a common user has a role in every container.

- **Local** — Roles created in a specific pluggable database (PDB). A local role exists only in a single pluggable database and can only contain roles and privileges that apply within the pluggable database in which the role exists.

Common role names must start with a `c##` prefix. Starting from Oracle 12.1.0.2, you can change these prefixes using the `COMMON_USER_PREFIX` parameter.

A `CONTAINER` clause can be added to `CREATE ROLE` statement to choose the container applicable for the role.

Examples

Create a common role.

```
show con_name

CON_NAME
CDB$ROOT

CREATE ROLE c##common_role;

Role created.
```

Create a local role.

```
show con_name

CON_NAME
ORCLPDB

CREATE ROLE local_role;

Role created.
```

Grant privileges and roles to the `local_role` database role.

```
GRANT RESOURCE, ALTER SYSTEM, SELECT ANY DICTIONARY TO local_role;
```

Database users to which the `local_role` role is granted now have all privileges that were granted to the role.

Revoke privileges and roles from the `local_role` database role.

```
REVOKE RESOURCE, ALTER SYSTEM, SELECT ANY DICTIONARY FROM local_role;
```

For more information, see [Overview of PL/SQL](#) in the *Oracle documentation*.

MySQL usage

Currently in MySQL 5.7, there is no ROLE feature. You must specify required privileges. However, there is an option when granting privileges to use wild-card characters to specify multiple privileges on one or more objects.

Note

Amazon Relational Database Service (Amazon RDS) for MySQL version 8 supports roles which are named collections of privileges. Roles can be created and dropped. Roles can have privileges granted to and revoked from them. Roles can be granted to and revoked from user accounts. The active applicable roles for an account can be selected from among those granted to the account and can be changed during sessions for that account.

For more information, see [Using Roles](#) in the *MySQL documentation*.

```
CREATE ROLE 'app_developer', 'app_read', 'app_write';
```

Note

Amazon RDS for MySQL version 8 incorporates the concept of user account categories with system and regular users distinguished according to whether they have the `SYSTEM_USER` privilege. For more information, see [Account Categories](#) in the *MySQL documentation*.

```
CREATE USER u1 IDENTIFIED BY 'password';  
  
GRANT ALL ON *.* TO u1 WITH GRANT OPTION;
```

```
-- GRANT ALL includes SYSTEM_USER, so at this point

-- u1 can manipulate system or regular accounts
```

Examples

Grant privileges using a wild-card.

```
GRANT ALL ON test_db.* to 'testuser';
GRANT CREATE USER on *.* to 'testuser';
GRANT SELECT ON db2.* TO 'testuser';
GRANT EXECUTE ON PROCEDURE mydb.myproc TO
```

For more information, see [GRANT Statement](#) in the *MySQL documentation*.

Oracle database users and MySQL users

With Amazon DMS, you can migrate existing Oracle and MySQL databases to Amazon Aurora with minimal downtime. Oracle Database users and MySQL users refer to the user accounts that have been granted specific privileges to access and manipulate data within the respective database systems. Scenarios where you might need to learn about these users include migrating production databases with existing user accounts, consolidating multiple databases with different user configurations, or recreating a development environment with the same user structure as production.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Syntax and option differences, similar functionality.

Oracle usage

Database user accounts are used for authenticating connecting sessions and authorizing access for individual users to specific database objects. Database Administrators grant privileges to user accounts, and applications use user accounts to access database objects.

Steps for providing database access to applications

1. Create a user account in the database. User accounts are typically authenticated using a password. Additional methods of authenticating users also exist.
2. Assign permissions to the database user account enabling access to certain database objects and system permissions.
3. Connecting applications, authenticate using the database username and password.

Oracle database users common properties

- Granting privileges or roles (collection of privileges) to the database user.
- Defining the default database tablespace for the user.
- Assigning tablespace quotas for the user.
- Configuring password policy, password complexity, lock, or unlock the account.

Authentication mechanisms

- **Username and password** — This is the default option.
- **External** — Using the operating system or third-party software, such as Kerberos.
- **Global** — Enterprise directory service, such as Active Directory or Oracle Internet Directory.

Oracle Schemas Compared to Users

In an Oracle database, a user equals a schema. This relationship is special because users and schemas are essentially the same thing. Consider an Oracle database user as the account you use to connect to a database while a database schema is the set of objects such as tables, views, and so on, that belong to that account.

- You can't create schemas and users separately. When you create a database user, you also create a database schema with the same name.
- When you run the `CREATE USER` command in Oracle, you create a user for login and a schema in which to store database objects.
- Newly created schemas are empty, but objects such as tables can be created within them.

Database users in Oracle 12c

Two types of users exist in the Oracle 12c database:

- **Common users** — Created in all database containers, root, and Pluggable Databases (PDB). Common users must have the C## prefix in the username.
- **Local users** — Created only in a specific PDB. Different database users with identical usernames can be created in multiple PDBs.

Examples

The following example demonstrates the following operations:

- Create a common database user using the default tablespace.
- Grant privileges and roles to the user.
- Assign a profile to the user, unlock the account, and force the user to change the password (PASSWORD EXPIRE).
- Create a local database user in the my_pdb1 pluggable database.

```
CREATE USER c##test_user IDENTIFIED BY password DEFAULT TABLESPACE USERS;
GRANT CREATE SESSION TO c##test_user;
GRANT RESOURCE TO c##test_user;
ALTER USER c##test_user ACCOUNT UNLOCK;
ALTER USER c##test_user PASSWORD EXPIRE;
ALTER USER c##test_user PROFILE ORA_STIG_PROFILE;
ALTER SESSION SET CONTAINER = my_pdb1;
CREATE USER app_user1 IDENTIFIED BY password DEFAULT TABLESPACE USERS;
```

For more information, see [Managing Security for Oracle Database Users](#) in the *Oracle documentation*.

MySQL usage

Database user accounts are used for authenticating connecting sessions and authorizing access for individual users to specific database objects. Database Administrators grant privileges to database user accounts that are used by applications to authenticate with an Aurora MySQL database.

For each account, `CREATE USER` creates a new row in the `mysql.user` system table. The account row reflects the properties specified in the statement. Unspecified properties are set to their default values:

- **Authentication** — The authentication plugin defined by the `default_authentication_plugin` system variable, and empty credentials.
- **SSL/TLS** — None.
- **Resource limits** — Unlimited.
- **Password management** — `PASSWORD EXPIRE DEFAULT`.
- **Account locking** — `ACCOUNT UNLOCK`.

When first created, accounts have no privileges. To assign privileges, use the `GRANT` statement.

Steps for providing database access to applications

1. Create a user account in the database. Typically, users authenticate using a username and password. Additional methods of authenticating users also exist.
2. Assign permissions to the database user account enabling access to certain database objects and system permissions.
3. Connecting applications, use the database username and password combination to authenticate with the database.

MySQL database users common properties

- Granting privileges to the database user.
- Configuring password policy, password complexity, lock, or unlock the account.
- Specifying authentication methods.
- User naming to indicate from which host names the user can login.
- Profiling, for example: `MAX_QUERIES_PER_HOUR` or `MAX_USER_CONNECTIONS`.

Authentication mechanisms

- **Username and password** — This is the default option.
- **External** — Using the operating system or third-party software, such as an IAM user.

- **Global** — Enterprise directory service, such as Active Directory.

IAM authentication

This feature is the equivalent to Oracle OS authentication.

With Amazon RDS for MySQL or Aurora MySQL, you can authenticate to your DB instance or DB cluster using Amazon Identity and Access Management (IAM) database authentication. With this authentication method, you don't need to use a password when you connect to a DB instance. Instead, you use an authentication token.

IAM database authentication provides the following benefits:

- Network traffic to and from the database is encrypted using Secure Sockets Layer (SSL).
- You can use IAM to centrally manage access to your database resources, instead of managing access individually on each DB instance or DB cluster.
- For applications running on Amazon EC2, you can use EC2 instance profile credentials to access the database instead of a password, for greater security.

Note

With IAM database authentication, you are limited to a maximum of 20 new connections in a single second.

Examples

The following example demonstrates the following operations:

- Create a database user using the `PASSWORD EXPIRE` option.
- Grant privileges to the user.
- Assign profiling properties to the user.

```
CREATE USER 'testuser'  
  IDENTIFIED BY 'new_password' PASSWORD EXPIRE;  
GRANT ALL ON test_db.* to 'testuser';  
GRANT CREATE USER on *.* to 'testuser';
```

```
ALTER USER 'testuser' WITH MAX_QUERIES_PER_HOUR 90;
```

To create an IAM user, make sure that the IAM user or role exists and is named by the same database username.

```
CREATE USER jane_doe IDENTIFIED WITH AWSAuthenticationPlugin AS 'RDS';
```

For more information, see [CREATE USER Statement](#) and [Specifying Account Names](#) in the *MySQL documentation* and [IAM database authentication for MariaDB, MySQL, and PostgreSQL](#) in the *Amazon Relational Database Service User Guide*.

Oracle and MySQL physical storage

This section includes pages related to Oracle and MySQL physical storage.

Topics

- [Oracle and MySQL table partitioning](#)
- [Oracle sharding](#)

Oracle and MySQL table partitioning

With Amazon DMS, you can implement table partitioning for Oracle and MySQL databases, which involves dividing a large table into multiple smaller partitions. Table partitioning helps manage and maintain large databases by improving query performance, facilitating data management operations, and reducing storage costs.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		Partitioning	Aurora MySQL doesn't support interval partitioning, partition advisor, preference partitioning, virtual column-based partitioning, and automatic list partitioning.

Oracle usage

The purpose of database partitioning is to provide support for very large tables and indexes by splitting them into smaller pieces. Each partition has its own name and definitions. They can be managed separately or collectively as one object. From an application perspective, partitions

are transparent. Partitioned tables behave the same as non-partitioned tables allowing your applications access using unmodified SQL statements. Table partitioning provides several benefits:

- **Performance improvements** — Table partitions help improve query performance by accessing a subset of a partition instead of scanning a larger set of data. Additional performance improvements can be achieved when using partitions and parallel query execution for DML and DDL operations.
- **Data management** — Table partitions facilitate easier data management operations (such as data migration), index management (creation, dropping, or rebuilding indexes), and backup/recovery. These operations are also referred to as Information Lifecycle Management (ILM) activities.
- **Maintenance operations** — Table partitions can significantly reduce downtime caused by table maintenance operations.

Oracle 18c introduces the following enhancements to partitioning.

- **Online Merging of Partitions and Subpartitions:** now it is possible to merge table partitions concurrently with Updates/Deletes and Inserts on a partitioned table.
- Oracle 18c also allows to modify partitioning strategy for the partitioned table: e.g. hash partitioning to range. This can be done both offline and online.

Oracle 19 introduces hybrid partitioned tables: partitions can now be both internal Oracle tables and external tables and sources. It is also possible to integrate both internal and external partitions together in a single partitioned table.

Hash table partitioning

When a partition key is specified (for example, a table column with a NUMBER data type), Oracle applies a hashing algorithm to evenly distribute the data (records) among all defined partitions. The partitions have approximately the same size.

The following example creates a hash partitioned table.

```
CREATE TABLE SYSTEM_LOGS
  (EVENT_NO NUMBER NOT NULL,
  EVENT_DATE DATE NOT NULL,
  EVENT_STR VARCHAR2(500),
  ERROR_CODE VARCHAR2(10))
```

```
PARTITION BY HASH (ERROR_CODE)
PARTITIONS 3
STORE IN (TB1, TB2, TB3);
```

List table partitioning

You can specify a list of discrete values for the table partitioning key in the description of each partition. This type of table partitioning enables control over partition organization using explicit values. For example, partition events by error code values.

The following example creates a list-partitioned table.

```
CREATE TABLE SYSTEM_LOGS
(EVENT_NO NUMBER NOT NULL,
EVENT_DATE DATE NOT NULL,
EVENT_STR VARCHAR2(500),
ERROR_CODE VARCHAR2(10))
PARTITION BY LIST (ERROR_CODE)
(PARTITION warning VALUES ('err1', 'err2', 'err3') TABLESPACE TB1,
PARTITION critical VALUES ('err4', 'err5', 'err6') TABLESPACE TB2);
```

Range table partitioning

Partition a table based on a range of values. The Oracle database assigns rows to table partitions based on column values falling within a given range. Range table partitioning is one of the most frequently used type of partitioning, primarily with date values. Range table partitioning can also be implemented with numeric ranges (1-10000, 10001- 20000...).

The following example creates a range-partitioned table.

```
CREATE TABLE SYSTEM_LOGS
(EVENT_NO NUMBER NOT NULL,
EVENT_DATE DATE NOT NULL,
EVENT_STR VARCHAR2(500))
PARTITION BY RANGE (EVENT_DATE)
(PARTITION EVENT_DATE VALUES
LESS THAN (TO_DATE('01/01/2015',
'DD/MM/YYYY')) TABLESPACE TB1,
PARTITION EVENT_DATE VALUES
LESS THAN (TO_DATE('01/01/2016',
'DD/MM/YYYY')) TABLESPACE TB2,
PARTITION EVENT_DATE VALUES
```

```
LESS THAN (TO_DATE('01/01/2017',
'DD/MM/YYYY')) TABLESPACE TB3);
```

Composite table partitioning

With composite partitioning, a table can be partitioned by one data distribution method, and then each partition can be further subdivided into sub-partitions using the same, or different, data distribution method(s). For example:

- Composite list-range partitioning.
- Composite list-list partitioning.
- Composite range-hash partitioning.

Partitioning extensions

Oracle provides additional partitioning strategies that enhance the capabilities of basic partitioning. These partitioning strategies include:

- Manageability extensions.
 - Interval partitioning.
 - Partition advisor.
- Partitioning key extensions.
 - Reference partitioning.
 - Virtual column-based partitioning.

Split partitions

You can use the `SPLIT PARTITION` statement to redistribute the contents of one partition, or sub-partition, into multiple partitions or sub-partitions.

```
ALTER TABLE SPLIT PARTITION p0 INTO
(PARTITION P01 VALUES LESS THAN (100), PARTITION p02);
```

Exchange partitions

The `EXCHANGE PARTITION` statement is useful to exchange table partitions in or out of a partitioned table.

```
ALTER TABLE orders EXCHANGE
PARTITION p_ord3 WITH TABLE orders_year_2016;
```

Subpartitioning tables

You can create subpartitions within partitions to further split the parent partition.

```
PARTITION BY RANGE(department_id)
SUBPARTITION BY HASH(last_name)
SUBPARTITION TEMPLATE
  (SUBPARTITION a TABLESPACE ts1,
   SUBPARTITION b TABLESPACE ts2,
   SUBPARTITION c TABLESPACE ts3,
   SUBPARTITION d TABLESPACE ts4)
(PARTITION p1 VALUES LESS THAN (1000),
 PARTITION p2 VALUES LESS THAN (2000),
 PARTITION p3 VALUES LESS THAN (MAXVALUE))
```

For more information, see [Partitioning Concepts](#) in the *Oracle documentation*.

Automatic list partitioning

Oracle 12c introduces automatic list partitioning. This enhancement enables automatic creation of new partitions for new values inserted into a list-partitioned table. An automatic list-partitioned table is created with only one partition. The database creates the additional table partitions automatically.

The following example creates an automatic list-partitioned table.

```
CREATE TABLE SYSTEM_LOGS
  (EVENT_NO NUMBER NOT NULL,
   EVENT_DATE DATE NOT NULL,
   EVENT_STR VARCHAR2(500),
   ERROR_CODE VARCHAR2(10))
PARTITION BY LIST (ERROR_CODE) AUTOMATIC
(PARTITION warning VALUES ('err1', 'err2', 'err3'))
```

For more information, see [Oracle Partitioning](#) in the *Oracle documentation*.

MySQL Usage

The table partitioning mechanism in MySQL is similar to Oracle and contains most of the Oracle table partitioning features. The only items not supported in MySQL table partitioning are the automatic features such as interval partitioning and automatic list partitioning. You can implement these features using triggers or procedures. For more information, see [Partitioning](#) in the *MySQL documentation*.

Note

Amazon Relational Database Service (Amazon RDS) for MySQL version 8 support the following partitioning options: `ADD PARTITION`, `DROP PARTITION`, `COALESCE PARTITION`, `REORGANIZE PARTITION`, and `REBUILD PARTITION ALTER TABLE`. You can use them with `ALGORITHM={COPY|INPLACE}` and `LOCK` clauses. `DROP PARTITION` with `ALGORITHM=INPLACE` deletes data stored in the partition and drops the partition. However, `DROP PARTITION` with `ALGORITHM=COPY` or `old_alter_table=ON` rebuilds the partitioned table and attempts to move data from the dropped partition to another partition with a compatible `PARTITION ... VALUES` definition. Data that cannot be moved to another partition is deleted.

MySQL basic table partitioning methods

Hash table partitioning

Partitioning by hash is used mostly to achieve an even distribution of data between the partitions. Make sure that you specify a column value or expression based on a column value to be hashed and the number of partitions into which the partitioned table is to be divided when creating the partitioned table.

Make sure that you use an SQL expression that returns an integer for the hash expression. The only permitted data types beside integer are date types and one of the following functions:

```
ABS, CEILING, DAY, DAYOFMONTH, DAYOFWEEK, DAYOFYEAR, DATEDIFF, EXTRACT, FLOOR, HOUR,
MICROSECOND, MINUTE, MOD, MONTH, QUARTER, SECOND, TIME_TO_SEC, TO_DAYS, TO_SECONDS,
UNIX_TIMESTAMP (with TIMESTAMP columns), WEEKDAY, YEAR, YEARWEEK
```

For other column types you can use `KEY` partitioning, which takes any column used as part or all of the table's primary key.

Examples

The following example creates a hash-partitioned table.

```
CREATE TABLE SYSTEM_LOGS
  (EVENT_NO INT NOT NULL,
  EVENT_DATE DATE NOT NULL,
  EVENT_STR VARCHAR(500),
  ERROR_CODE INT)
PARTITION BY HASH (ERROR_CODE)
PARTITIONS 3;
```

The following example creates a key-partitioned table.

```
CREATE TABLE SYSTEM_LOGS
  (EVENT_NO INT NOT NULL,
  EVENT_DATE DATE NOT NULL,
  EVENT_STR VARCHAR(500),
  ERROR_CODE VARCHAR(10) PRIMARY KEY)
PARTITION BY KEY()
PARTITIONS 3;
```

For more information, see [HASH Partitioning](#) and [KEY Partitioning](#) in the *MySQL documentation*.

List table partitioning

As with the hash partition, make sure that this partitioned column is INT. To use LIST on `varchar`, use `LIST COLUMNS`.

Examples

The following example creates a list-partitioned table.

```
CREATE TABLE SYSTEM_LOGS
  (EVENT_NO INT NOT NULL,
  EVENT_DATE DATE NOT NULL,
  EVENT_STR VARCHAR(500),
  ERROR_CODE INT)
PARTITION BY LIST (ERROR_CODE)
(PARTITION warning VALUES IN (3345, 5423,3332),
PARTITION critical VALUES IN (9786, 9231, 6321));
```

The following example creates a list-columns-partition table.

```
CREATE TABLE SYSTEM_LOGS
  (EVENT_NO INT NOT NULL,
  EVENT_DATE DATE NOT NULL,
  EVENT_STR VARCHAR(500),
  ERROR_CODE VARCHAR(500))
PARTITION BY LIST COLUMNS (ERROR_CODE)
(PARTITION warning VALUES IN ('err1', 'err2', 'err3'),
PARTITION critical VALUES IN ('err4', 'err5', 'err6'));
```

For more information, see [LIST Partitioning](#) in the *MySQL documentation*.

Range table partitioning

Similar to a list partition, you can use a range partition on integer values or RANGE COLUMNS for DATE or DATETIME.

Examples

The following example creates a range-partitioned table.

```
CREATE TABLE SYSTEM_LOGS
  (EVENT_NO INT NOT NULL,
  EVENT_DATE DATE NOT NULL,
  EVENT_STR VARCHAR(500))
PARTITION BY RANGE (YEAR(EVENT_DATE))
(PARTITION p0 VALUES LESS THAN (2015),
PARTITION p1 VALUES LESS THAN (2016),
PARTITION p2 VALUES LESS THAN (2017));
```

The following example creates a range columns-partitioned table.

```
CREATE TABLE SYSTEM_LOGS
  (EVENT_NO INT NOT NULL,
  EVENT_DATE DATE NOT NULL,
  EVENT_STR VARCHAR(500))
PARTITION BY RANGE COLUMNS (EVENT_DATE)
(PARTITION p0 VALUES LESS THAN ('2015-01-01'),
PARTITION p1 VALUES LESS THAN ('2016-01-01'),
PARTITION p2 VALUES LESS THAN ('2017-01-01'));
```

For more information, see [RANGE Partitioning](#) in the *MySQL documentation*.

Composite table partitioning

With composite partitioning, a table can be partitioned by one data distribution method, and then each partition can be further subdivided into sub-partitions using the same, or different, data distribution methods.

In MySQL 5.7, you can subpartition tables that are partitioned by range or list. Subpartitions may use either hash or key partitioning.

You can use the following approaches:

- Specify only the number of subpartitions for each partition.
- Explicitly define subpartitions in any partition individually, this option is useful if you want to pick the names for your subpartitions.

Note

Make sure that all partitions have the same number of subpartitions.

Examples

The following example creates a range-key subpartition. All partitions have two subpartitions.

```
CREATE TABLE EMPLOYESS
  (DEPARTMENT_ID INT NOT NULL,
  LAST_NAME VARCHAR(50) NOT NULL,
  FIRST_NAME VARCHAR(50),
  PRIMARY KEY (DEPARTMENT_ID, LAST_NAME))
PARTITION BY RANGE(DEPARTMENT_ID)
SUBPARTITION BY KEY (last_name)
SUBPARTITIONS 2
  (PARTITION p1 VALUES LESS THAN (10),
  PARTITION p2 VALUES LESS THAN (20),
  PARTITION p3 VALUES LESS THAN (MAXVALUE));
```

For more information, see [Subpartitioning](#) in the *MySQL documentation*.

Split partitions

In Oracle, `SPLIT PARTITION STATEMENT` translates to `REORGANIZE PARTITION` in MySQL. Create a list partition and then split one of the partitions.

You can split range partitions at the last partition only.

```
CREATE TABLE SYSTEM_LOGS
  (EVENT_NO INT NOT NULL,
  EVENT_DATE DATE NOT NULL,
  EVENT_STR VARCHAR(500),
  ERROR_CODE VARCHAR(500))
  PARTITION BY LIST COLUMNS (ERROR_CODE)
  (PARTITION warning VALUES IN ('err1', 'err2', 'err3'),
  PARTITION critical VALUES IN ('err4', 'err5', 'err6'));

ALTER TABLE SYSTEM_LOGS REORGANIZE PARTITION warning INTO
  (PARTITION warning0 VALUES IN ('err2.5', 'err3.5'),
  PARTITION warning1 VALUES IN ('err2.8', 'err3.8'));
```

For more information, see [Management of RANGE and LIST Partitions](#) in the *MySQL documentation*.

Exchange partitions

Similar to Oracle, you can exchange tables with partitions.

```
ALTER TABLE orders
  EXCHANGE PARTITION p_ord3 WITH TABLE orders_year_2016;
```

For more information, see [Exchanging Partitions and Subpartitions with Tables](#) in the *MySQL documentation*.

Summary

Oracle table partition type	Built-in MySQL support	Example
List	Yes	LIST Partitioning
Range	Yes	RANGE Partitioning

Oracle table partition type	Built-in MySQL support	Example
Hash	Yes	HASH Partitioning
Composite or subpartitioning	Yes	Subpartitioning
Interval	No	Restrictions and Limitations on Partitioning
Partition advisor	No	Restrictions and Limitations on Partitioning
Preference	No	Restrictions and Limitations on Partitioning
Virtual column-based	No	Restrictions and Limitations on Partitioning
MySQL partitioning automatic list partitioning	No	Restrictions and Limitations on Partitioning
Split and exchange	Yes	ALTER TABLE Partition Operations and Exchanging Partitions and Subpartitions with Tables

Oracle sharding

With Amazon DMS, you can shard an Oracle database across multiple data stores to scale write throughput and distribute data horizontally. Oracle sharding is a database architecture that partitions data across multiple Oracle databases, providing horizontal scalability and improved performance.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
		N/A	MySQL doesn't support sharding.

Oracle usage

Sharding is a method of data architecture where table data is horizontally partitioned across independent databases. These databases are called shards. All of the shards make up a single logical database, which is referred to as a sharded database (SDB). Sharding a table is process of splitting this table between different shards where each shards will have sharded table with the same structure but different subset of rows.

Oracle 18c introduces following sharding enhancements:

- User-defined sharding. Before Oracle 18c data was redirected across shards by system. With user-defined sharding, users are now able to explicitly redirect sharded table data to specific individual shards.
- Using JSON, BLOB, CLOB and spatial objects functionality in a sharded environment. You can now use these objects in sharded tables.

For more information, see [Overview of Oracle Sharding](#) in the *Oracle documentation*.

MySQL usage

There is no equivalent option in MySQL. The most equivalent option will be to create application level sharding management that will interact with data that is spread across multiple instances.

Another option will be to assess the requirements and probably use another data store such as Amazon Redshift, Amazon EMR, or Amazon DynamoDB.

Oracle and MySQL monitoring

This section provides information about Oracle and MySQL monitoring.

Feature compatibility	Amazon SCT / Amazon DMS automation level	Amazon SCT action code index	Key differences
	N/A	N/A	Make sure to change table names in queries when using MySQL.

Oracle usage

Oracle provides several built-in views used to monitor the database and query its operational state. You can use these views to track the status of the database, view information about database schema objects, and obtain other information.

The data dictionary is a collection of internal tables and views that supply information about the state and operations of an Oracle database including database status, database schema objects such as tables, views, sequences, and so on, users and security, and physical database structure (datafiles). The contents of the data dictionary are persisted to disk.

Examples of data dictionary views include:

- `DBA_TABLES` — Information about all tables in the current database.
- `DBA_USERS` — Information about all database users.
- `DBA_DATA_FILES` — Information about all physical data files in the database.
- `DBA_TABLESPACES` — Information about all tablespaces in the database.
- `DBA_TAB_COLS` — Information about columns (for all tables) in the database.

Note

Data dictionary view names can start with `DBA_*`, `ALL_*`, `USER_*`, depending on the presented level and scope of information (user-level or database-level).

For more information, see [Static Data Dictionary Views](#) in the *Oracle documentation*.

Dynamic performance views (V\$ Views) are a collection of views that provide real-time monitoring information about the current state of the database instance configuration, runtime statistics, and operations. These views are continuously updated while the database is running.

Information provided by the dynamic performance views includes session information, memory usage, progress of jobs and tasks, SQL execution state and statistics, and various other metrics.

Common dynamic performance views include:

- `V$SESSION` — Information about all current connected sessions in the instance.
- `V$LOCKED_OBJECT` — Information about all objects in the instance on which active locks exist.
- `V$INSTANCE` — Dynamic instance properties.
- `V$SESSION_LONG_OPS` — Information about certain long-running operations in the database such as queries currently executing.
- `V$MEMORY_TARGET_ADVICE` — Advisory view on how to size the instance memory, based on instance activity and past workloads.

For more information, see [Data Dictionary and Dynamic Performance Views](#) in the *Oracle documentation*.

MySQL usage

MySQL provides two different ways to retrieve information about the state of the database and current activities. The information is similar in nature to the Oracle data dictionary tables and V\$ performance views. In addition, Amazon Aurora MySQL provides a Performance insights console for monitoring and analyzing database workloads and troubleshooting performance issues.

Information schema tables

The information schema consists of views containing information about objects that were created in the current database.

- The information schema is specified by the SQL standard and is supported by MySQL.
- Some of these tables are comparable to Oracle `USER_*` Data Dictionary tables.
- The owner of this schema is the initial database user.
- Since the information schema is defined as part of the SQL standard, it can be expected to remain stable across MySQL versions.

SHOW command

The `SHOW` command provides information about databases, tables, columns, and status information about the server.

- If the syntax for a `SHOW` statement includes a `LIKE` pattern part, the pattern is a string that can contain the SQL `%` and `_` wildcard characters. The pattern is useful for restricting statement output to matching values.
- The `SHOW` command has more dynamic views such as `PROCESSLIST`.
- Users must have `PROCESS` privilege to query this data.

Summary

Information	Oracle	MySQL
Database properties	<code>V\$DATABASE</code>	<code>pg_database</code>
Database sessions	<code>V\$SESSION</code>	<code>SHOW PROCESSLIST</code>
Database users	<code>DBA_USERS</code>	<code>mysql.user</code>
Database tables	<code>DBA_TABLES</code>	<code>information_schema .TABLES</code>

Information	Oracle	MySQL
Database data files	DBA_DATA_FILES	information_schema .FILES
Table columns	DBA_TAB_COLS	information_schema .COLUMNS
Database locks	V\$LOCKED_OBJECT	information_schema .INNODB_LOCKS
Currently configured runtime parameters	V\$PARAMETER	SHOW GLOBAL VARIABLES
All system statistics	V\$SYSSTAT	information_schema .INNODB_METRICS
Privileges on tables	DBA_TAB_PRIVS	information_schema .TABLE_PRIVILEGES
Information about IO operations	V\$SEGSTAT	SHOW STATUS LIKE '%read%'; SHOW STATUS LIKE '%write%';

For more information, see [SHOW Statements](#) and [INFORMATION_SCHEMA Tables](#) in the *MySQL documentation*.

Oracle to [.shared] `AURMySQL` migration quick tips

This section provides migration tips that can help save time as you transition from Oracle to Amazon Aurora MySQL-Compatible Edition (Aurora MySQL). These tips address many of the challenges faced by administrators new to Aurora MySQL. Some of these tips describe functional differences in similar features between Oracle and Aurora MySQL.

Management

- In Aurora MySQL, *database snapshot* is equivalent to RMAN backup in Oracle.
- Partitioning in Aurora MySQL doesn't provide many of the Oracle features such as Partition Advisor, Preference Partitioning, Virtual Column-Based Partitioning, and Automatic List Partitioning.
- Unlike Oracle statistics, Aurora MySQL doesn't collect detailed key value distribution in tables. Aurora MySQL only collects statistics on indexes.
- You can use Amazon services, such as Lambda, to replicate functionality of features not provided by MySQL, such as email.
- Amazon RDS manages parameters and backups. It is very useful for checking a parameter's value against its default and comparing them to another parameter group.
- With just a few clicks, you can create replicas to implement high availability.
- Aurora MySQL doesn't have an equivalent to database links. Aurora MySQL can only query across databases within the same instance.

SQL

- Aurora MySQL doesn't support statement-level triggers or triggers on system events.
- Aurora MySQL doesn't support many cursor status checks. When you declare cursors in Aurora MySQL, make sure that you create an explicit HANDLER object.
- To run a stored procedure or function, use CALL instead of EXECUTE.
- To run a string as a query, use Aurora MySQL Prepared Statements instead of EXECUTE(<String>).
- In Aurora MySQL, make sure that you terminate the IF blocks with END IF. Also, make sure that you terminate the WHILE . . LOOP loops with END LOOP.

- Unlike Oracle, Aurora MySQL auto-commit default is set to ON. Be sure to set it to OFF to enable the database behavior similar to Oracle.
- Similar to Oracle, you can define collations at the server, database, and column level. You can't define collations at the table level in Aurora MySQL.
- In Oracle, the DELETE <Table Name> syntax enables you to omit the FROM keyword. This syntax is not valid in Aurora MySQL. Add the FROM keyword to all DELETE statements.
- In Aurora MySQL, the AUTO_INCREMENT column property is similar to IDENTITY in Oracle.
- Error handling in Aurora MySQL has less features than Oracle. For special requirements, you can log or send alerts by inserting into tables or catching errors.
- Aurora MySQL doesn't support the MERGE statement. Use the REPLACE statement and the INSERT...ON DUPLICATE KEY UPDATE statement as alternatives.
- Unlike Oracle, you can't concatenate strings in Aurora MySQL using the || operator.
- Aurora MySQL is much stricter than Oracle for statement terminators. Make sure that you use semicolons at the end of statements.
- Aurora MySQL doesn't support the BFILE, ROWID, and UROWID data types.
- In MySQL, temporary tables are retained only for the session and only the session that created a temporary table can query it.
- MySQL doesn't support unused or virtual columns and there is no workaround for replacing unused columns to achieve functionality similar to virtual columns. You can combine views and functions.
- MySQL doesn't support materialized views. Use views or summary tables instead.
- Explore Amazon to locate features that can be replaced with Amazon services. They can help you maintain your database and decrease costs.
- In MySQL, you can create multiple databases in a single instance. This approach can be useful for consolidation projects.
- Beware of control characters when copying and pasting a script to Aurora MySQL clients. Aurora MySQL is much more sensitive to control characters than Oracle and can result in frustrating syntax errors that are hard to find.