ORACLE

OCI Database Migration service End-To-End AWS RDS to ADB-S Online Migration Tutorial

Aimed for scenarios where your application must remain online, and your source database has a direct connection to OCI.

May 2025, Version 2.0
Copyright © 2025, Oracle and/or its affiliates
Public

Purpose statement

This document walks you through all the steps required to migrate an existing AWS RDS for Oracle database to an OCI Autonomous database (ADB). You will provision a Virtual Cloud Network (VCN) and an ADB instance to be used as a target then you will perform a database migration using OCI Database Migration service (DMS).

With DMS we make it quick and easy for you to migrate databases from on-premises, Oracle Cloud Infrastructure, or third-party cloud into Oracle databases on OCI.

Disclaimer

This document in any form, software, or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced, or distributed to anyone outside Oracle without the prior written consent of Oracle. This document is not part of your license agreement, nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remain at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

Table of contents

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

Introduction to OCI Database Migration – DMS

OCI Database Migration (DMS) provides a high performant, self-service experience to achieve migrations, which include:

Homogeneous migration of data from MySQL or Oracle databases into OCI.

Provides enterprise-level logical Online and Offline migrations with minimal downtime based on industry leading GoldenGate for data replication.

DMS Documentation:

Please review the documentation here.

Assumptions

- There is an existing AWS RDS for Oracle instance with Oracle Enterprise Edition, version 19 was selected for this exercise. The following link contains more information on how to create one.
- Database archiving mode set to ARCHIVELOG, this happens automatically when automated backups are enabled by setting the backup retention period to a value greater than 0.
- Set the instance as Publicly accessible.
- Selected architecture for this exercise is Non-CDB. Learn more about RDS for Oracle architecture at the following link.
- Amazon S3 integration is enabled, the database can transfer files between RDS for Oracle DB instance and an Amazon S3 bucket. To learn more, check the following link.
- The connection to the target database is done using cloud shell, the networking needs to match the one in the target database. Since cloud shell will only list VCNs from the tenancy home region this is the region that will be selected for the migration.

Task 0: Understand New DMS Concepts

DMS provides a fully managed approach to migrating databases from various locations into OCI-hosted databases.

Migrations can be either one of the following modes:

- **Offline**: The Migration makes a point-in-time copy of the source to the target database. Any changes to the source database during migration are not copied, requiring any applications to stay **offline** for the duration of the migration.
- **Online**: The Migration makes a point-in-time copy and replicates all subsequent changes from the source to the target database. This allows applications to stay **online** during the migration and then be switched over from source to target database.

DMS supports both offline and online mode. The service supports Oracle databases located on-premises, in 3rd party clouds, or on Oracle OCI as source. The targets can be Oracle Autonomous Database shared or dedicated and Oracle Cloud Infrastructure co-managed Oracle Base Database service (Oracle Base Database (VM, BM) and Exadata on Oracle Public Cloud).

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

The DMS service runs as a managed cloud service separate from the user's tenancy and resources. The service operates as a multi-tenant service in a DMS Service Tenancy and communicates with the user's resources using Private Endpoints (PEs). PEs are managed by DMS and are transparent to the user.

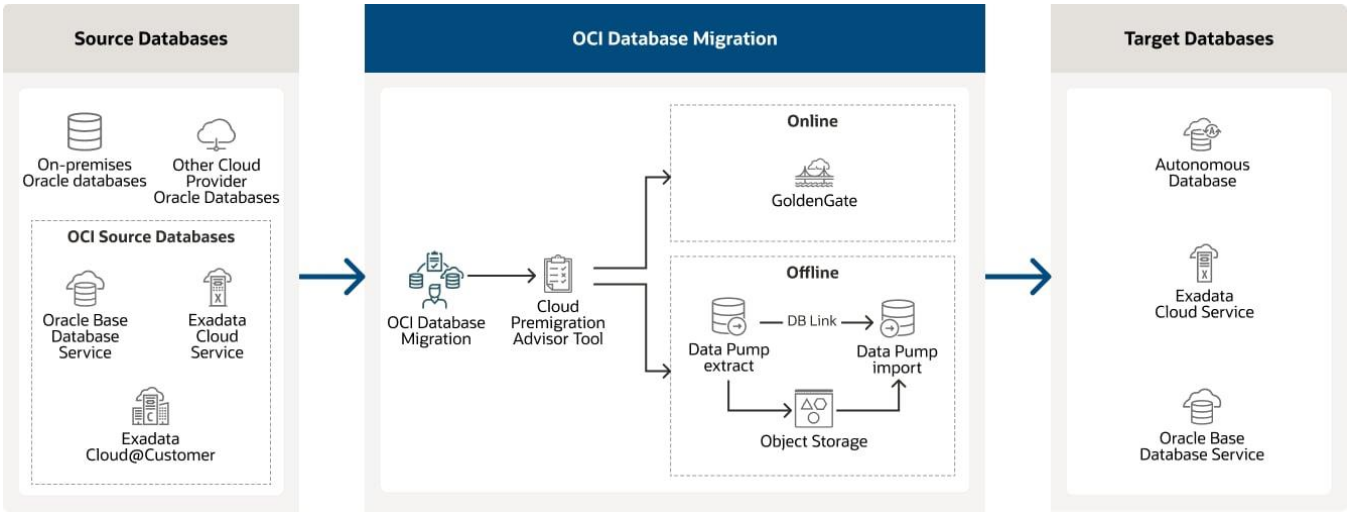

DMS Simplified Topology

Compartment: A compartment is a collection of related resources (such as cloud networks, compute instances, or block volumes) that can be accessed only by those groups that have been given permission by an administrator in your organization. For example, one compartment could contain all the servers and storage volumes that make up the production version of your company's Human Resources system. Only users with permission to that compartment can manage those servers and volumes.

Data region: A geographical region that's associated with one or more data centers. When you sign up for an Oracle Cloud account, you select a default data region, where your services will be hosted.

DMS Control Plane: Used by DMS end user to manage Migration and database connection objects. The control plane is exposed through the DMS Console UI as well as the Rest API.

DMS Data Plane: Managed by DMS Control Plane and transparent to the user. The GGS Data Plane manages ongoing migration jobs and communicates with the user's databases and GoldenGate instance using PEs. The DMS data plane does not store any customer data, as data flows through GoldenGate and Data Pump directly within the user's tenancy.

Migration: A Migration contains metadata for migrating one database. It contains information about source, target, and migration methods and is the central object for users to run migrations. After creating a migration, a user can validate the correctness of the environment and then run the migration to perform the copy of database data and schema metadata from source to target.

Migration Job: A Migration Job displays the state or a given Migration execution, either for validation or migration purposes. A job consists of several sequential phases, users can opt to wait after a given phase for user input to resume with the following phase.

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

Database connection: A database connection represents information about a source or target database, such as connection and authentication credentials. DMS uses the OCI Vault to store credentials. A database connection is reusable across multiple migrations.

Task 1: Have the Administrator Set Required Permissions

The following permissions need to be set to have access to the necessary objects unless you have administrative privileges. The following permissions assume that the user is part of group DMS_LA and all resources are created in a compartment called DMS_LA. Have your tenancy administrator set these permissions.

PERMISSIONS REQUIRED BY DMS TO USE DATABASES, VAULTS, AND NETWORKING
Allow group DMS_LA to manage virtual-network-family in compartment DMS_LA
Allow group DMS_LA to manage vaults in compartment DMS_LA
Allow group DMS_LA to manage keys in compartment DMS_LA
Allow group DMS_LA to manage database-family in compartment DMS_LA
Allow group DMS_LA to manage autonomous-database-family in compartment DMS_LA
Allow group DMS_LA to manage object-family in compartment DMS_LA
Allow group DMS_LA to manage secret-family in compartment DMS_LA
Allow group DMS_LA to manage goldengate-connections in compartment DMS_LA
Allow group DMS_LA to manage odms-connection in compartment DMS_LA
Allow group DMS_LA to manage odms-migration in compartment DMS_LA
Allow group DMS_LA to manage odms-job in compartment DMS_LA
Allow group DMS_LA to manage cloud-shell in compartment DMS_LA

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

Task 2: Sign in and Open DMS Console

To perform this guide, you need to have access to an OCI tenancy with access to a region where DMS is released, such as the US-Ashburn-1 region. Please review https://www.oracle.com/cloud/data-regions/ for available regions.

1. Open the browser with URL https://console.us-ashburn-1.oraclecloud.com/ (Adjust for home region)
2. Log in using your tenancy name and username/password.
3. In the OCI console title bar change region if applicable.

Task 3: Create Virtual Cloud Network

The following task is optional if a suitable VCN is already present.

1. In the OCI Console Menu, go to Networking > Virtual Cloud Networks
2. Pick a compartment on the left-hand side Compartment list. You need to have the necessary permissions for the compartment.
3. Press Actions >Start VCN Wizard and pick VCN with Internet Connectivity.
4. Enter a VCN Name, such as VCN_DMS_LA. Leave CIDR block defaults, unless you need non-overlapping addresses for peering later. Press Next.
5. Review Summary and press Create.

Task 4: Update Security List for Virtual Cloud Network Subnet

This task assumes default permissions in your public subnet. If you disabled or restricted your default permissions such as port 22 SSH access or restricted egress, please add default permissions as needed.

1. In the OCI Console Menu, go to **Networking > Virtual Cloud Networks** and pick your VCN.
2. Navigate to the Subnets tab, pick Public Subnet- *VCN NAME*.
3. Navigate to Security tab, in the Security Lists list pick Default Security List for *VCN NAME*.
4. Navigate to Security rules tab, in the Ingress Rules list press Add Ingress Rules.
5. Enter the following values, otherwise leave defaults:
    a. Source CIDR: 0.0.0.0/0
    b. Destination Port Range: 443
    c. Description: OGG HTTPS
6. Close dialog by pressing **Add Ingress Rules**.
7. In the Ingress Rules list press Add Ingress Rules.
8. Enter the following values, otherwise leave defaults:
    a. Source CIDR: **10.0.0.0/16**
    b. Destination Port Range: **1521**
    c. Description: Oracle DB access for PEs
9. Close dialog by pressing **Add Ingress Rules**.

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

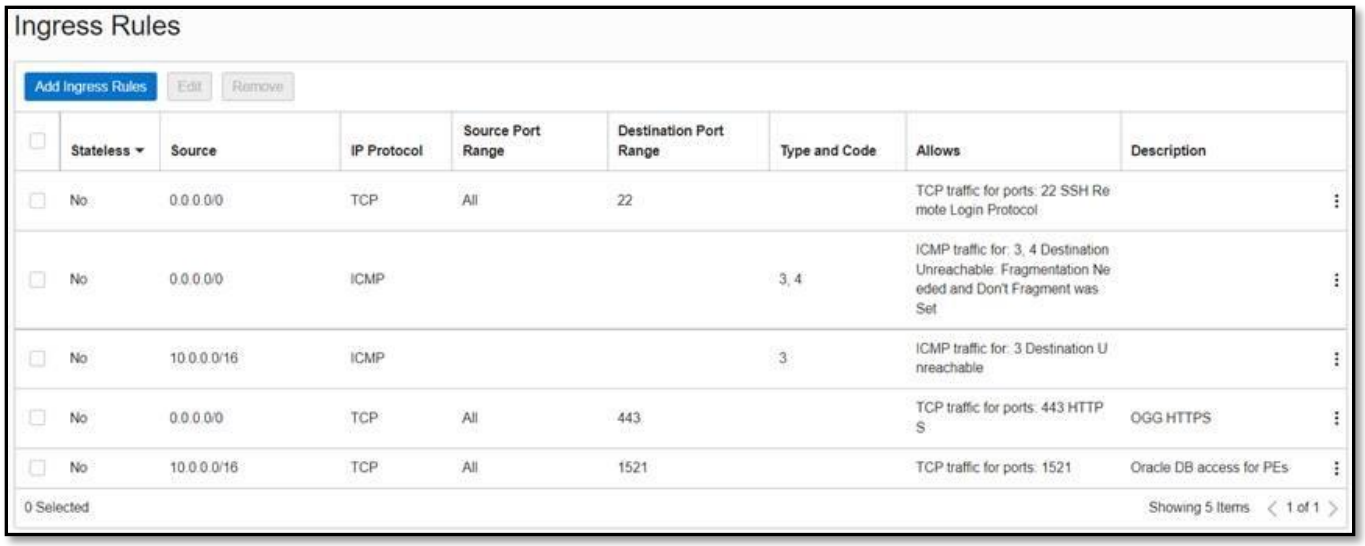

Task 5: Create Vault

The following task is optional if a Vault is already present.

1. In the OCI Console Menu, go to **Identity & Security > Vault**.
2. Pick a compartment on the left-hand side **Compartment** list.
3. Press Create Vault.
4. In the **Create Vault** dialog, enter a Name such as **DMS_Vault**.
5. Close the dialog by pressing **Create Vault**.
6. Wait until the state of the new vault is **Active**.
7. Click on the new vault and press **Create Key** in the **Master Encryption Keys** list.
8. In the **Create Key** dialog, enter a Name such as **DMS_Key**.
9. Close the dialog by pressing **Create Key**.

Task 6: Identify the RDS instance details

First find the endpoint (DNS name) and port number for the RDS DB instance.
Navigation: Amazon RDS homepage>Databases>Your DB> **Connectivity & security tab**:

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

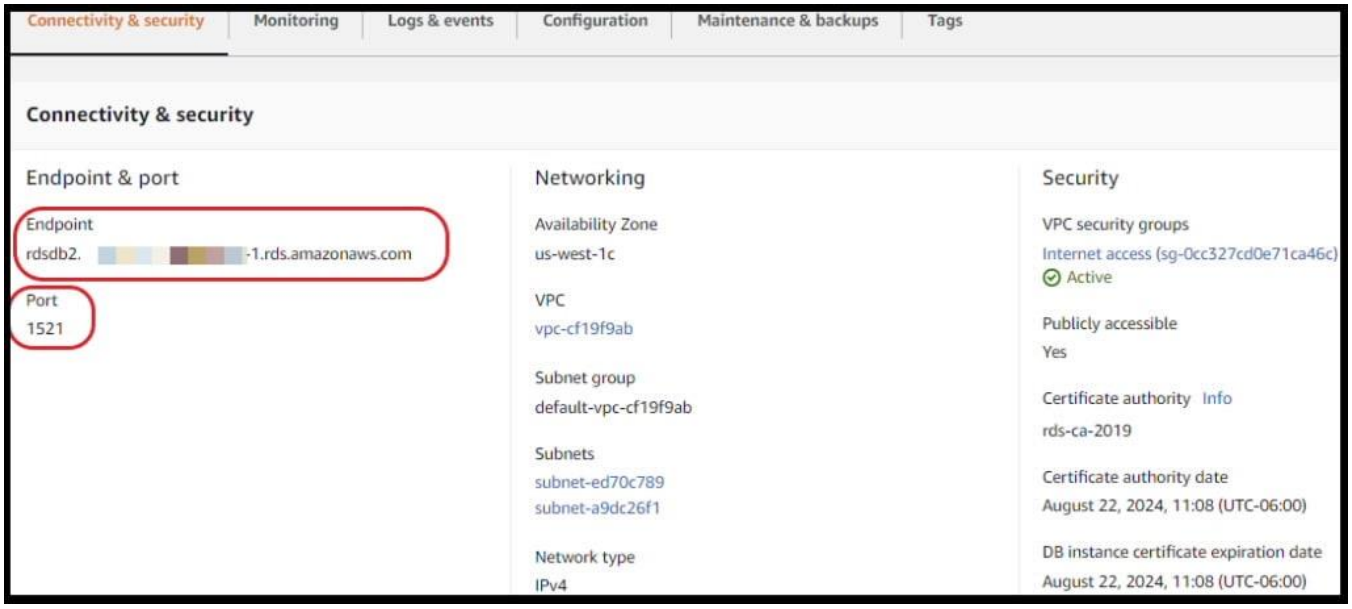

* OCI Database Migration needs an IP address, the following command nslookup + RDS Private Endpoint should show a similar response as the following, take note of the IP address:

```
nslookup  .us-east-1.rds.amazonaws.com

Name:  .us-east-1.rds.amazonaws.com
Address 1: 44. .152.184 ec2-44- -152-184.compute-1.amazonaws.com
```

In the **Configuration** tab, locate the following information:

- **DB name** (not the DB instance ID)
- **Master username**

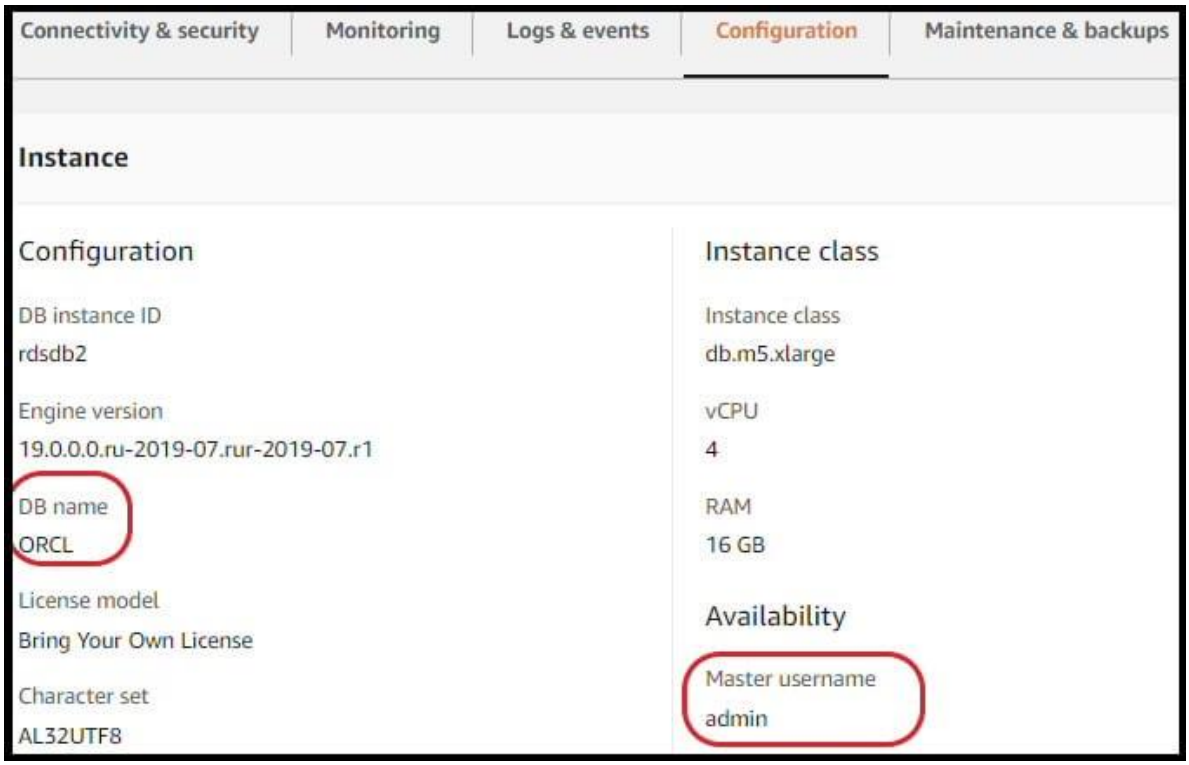

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

Navigate to the AWS Console Home>S3>**buckets**:

Identify the bucket **Name** and **Region**.

Task 7: Create Target Autonomous Database

The following task is optional if a target autonomous database is already present. In this example the target database is an ATP-shared instance with private IP address.

1. You first need to create a Network Security Group for use in a Private IP ADB instance. In the OCI Console Menu, go to **Networking > Virtual Cloud Networks** and pick your VCN.
2. Navigate to Security tab, **Network Security Groups**.
3. Press **Create Network Security Group**.
4. Enter Name such as **DMS_NSG** and press **Next.**
5. In the **Rule** box please enter the following entries, otherwise leave defaults**:**

   Source Type: CIDR

   Source CIDR: 0.0.0.0/0
6. Press **Create**.
7. Now you can create the ADB instance. In the OCI Console Menu, go to **Oracle Database > Autonomous Database >Autonomous Transaction Processing.**
8. Pick a compartment on the Applied filters.
9. Press **Create Autonomous Database**.
10. Enter the following values, otherwise leave defaults. You can adjust shapes and storage to your use case.
    - Display Name: TargetATP
    - Database name: TargetATP
    - Workload type: Transaction Processing
    - Advanced options> Bring Your Own License (BYOL): Enable
    - Create administrator credentials – Password: password of your choice
    - Network access > Access Type: Private endpoint access only
    - Virtual cloud network: VCN_DMS_LA (Or your VCN name)
    - Client subnet: Private Subnet-VCN_DMS_LA (Or your subnet name)
    - Advanced options> Network security group: DMS_NGS (Or your NSG name)
11. Close the dialog by pressing **Create Autonomous Database.**

Task 9: Prepare the AWS RDS Oracle database

This task prepares required user accounts and settings for migration in the source database. It assumes default settings in the database. If you changed the default settings, further settings might be necessary.

1) Download the preparation script from this link

2) Locate the file and run it ./dms-db-prep-v2.sh

3) Follow the instructions:

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

a) Database type [(s)ource/(t)arget]?: s

b) Is your source database hosted in AWS RDS (Amazon Relational Database Service)? [y/n]: y

c) Migration type [(on)line/(off)line]: on

d) Are you using ggadmin as the replication database username (recommended)? [y/n]: y

e) Password for ggadmin user:

```
-- Oracle Cloud Infrastructure Database Migration Service --
This script will help you prepare your source and target databases for migration.
Please answer the following questions to proceed:

Database type [(s)ource/(t)arget]?: s
Is your source database hosted in AWS RDS (Amazon Relational Database Service)? [y/n]: y
Migration type [(on)line/(off)line]: on
Are you using ggadmin as the replication database username (recommended)? [y/n]: y
Password for ggadmin user:

Sql script /home/mobaxterm/desktop/dms_prep_db.sql generated.
Please connect to your database as sysdba (role) and run the above generated sql script.
This script will analyze your database and will generate a subsequent sql script that you must review, modify (if needed) and run in order to get your database set up for the migration
[    info] In order to have your database prepared for the migration, please set the following parameters through the Parameter groups functionality:
[    info] STREAMS_POOL_SIZE=2147483648
[    info] ENABLE_GOLDENGATE_REPLICATION=TRUE
[    info] GLOBAL_NAMES=FALSE
[    info] To see how Parameter groups work please refer to https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/parameter-groups-overview.html
[    info]
[    info] When setting up your migration through the OCI Console:
[    info] Use ggadmin user as the GoldenGate Database User for Source Database when creating the Migration definition through the Database Migration Service
```

The following parameters need to be set with the **Parameter groups** functionality:

- STREAMS_POOL_SIZE=2147483648
- ENABLE_GOLDENGATE_REPLICATION=TRUE
- GLOBAL_NAMES=FALSE

To see how Parameter groups work please refer to the following link.

The next step is to locate the output file dms_prep_db.sql generated by the script, you should see the path on the screen. Connect to your database as sysdba (role) and run the above generated sql script.

This script will analyze your database and will generate a subsequent sql script (DMS_Configuration.sql) that you must review, modify (if needed) and run in order to get your database set up for the migration.

```
  SELECT program INTO :NEW.session_program FROM v$session a
    WHERE a.SID = sys_context('USERENV','SID');
  EXCEPTION
    WHEN OTHERS THEN
      v_code := SQLCODE;
      v_errm := SUBSTR(SQLERRM, 1, 64);
      DBMS_OUTPUT.PUT_LINE('Error on GGADMIN.TRG_INSERT_EVENT_TABLE');
      DBMS_OUTPUT.PUT_LINE(v_code || ': ' || v_errm);
      RAISE;
END;
--
-- Script DMS_Configuration.sql generated. Please review this script, modify as appropriate and run it in your database.
-- Your source database will be ready for migration after execution of these operations.
```

Task 10: Prepare the target Autonomous Database

The next steps will connect to the target ADB instance and enable the standard ggadmin user, you can skip these steps if the user is already enabled.

Make sure that your Autonomous Database mTLS authentication option is marked as 'Not required', you can check this in the following navigation path: Overview/Autonomous Database/Autonomous Database details

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

Go to Database connection/ Connection settings section and select TLS from the TLS authentication list of values, then copy the connection string for one of the TNS names.

Click Download wallet from the section “Download client credentials (Wallet)”

Open cloud shell in OCI, this icon is available in the top ribbon, next to the region:

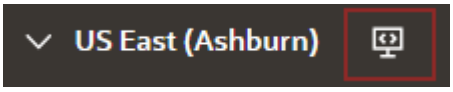

In the Cloud Shell, click on the Upload button (it looks like a cloud with an up arrow).

Upload the wallet zip file you downloaded.

In the Cloud Shell, unzip the wallet file using the following command:

```
unzip Wallet_<your_db_name>.zip -d wallet
```

Set the TNS_ADMIN environment variable to point to the directory where you unzipped the wallet, i.e:

```
export TNS_ADMIN=$HOME/wallet
```

One of the extracted files from the wallet is sqlnet.ora, make sure that in the content the directory path is the same path where you decompressed the wallet file, i.e:

```
WALLET_LOCATION = (SOURCE = (METHOD = file) (METHOD_DATA =
(DIRECTORY="/home/cloudshell-user/wallet")))
```

Connect to sqlplus using the next command, get the connection string from your database “Connection strings” section:

```
sqlplus admin/ <ADB password>@ADB connection string
```

In SQL Plus enter the following commands:

```
SQL> alter user ggadmin identified by <new password> account unlock;
User altered.
SQL> quit
```

Task 11: Create Object Store Bucket for Data Pump Storage

Object Store is used as temporary storage between source and target databases with Data Pump. This task is creating an empty bucket for use in the migration.

In the OCI Console Menu, go to Storage > Object Storage & Archive…

Press Create Bucket.

On the page Create Bucket, fill in the following entries, otherwise leave defaults:

- Bucket Name: **DMSStorage**

Press Create Bucket

Task 12: Create a database connection for the RDS source database

Database connection resources enable networking and connectivity for the source and target databases.

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

In the OCI Console Menu, go to Migration & Disaster Recovery> Database Migration > Database Connections.

Press Create connection.

In the page Database Details, fill in the following entries, otherwise leave defaults:

- Name: **RDS**
- Type: **Amazon RDS for Oracle**
- Vault: **DMS_Vault**
- Encryption Key: **DMS_Key**
- Database connection string: Provide the RDS public IP, port, and the database name, i.e:
    a. 44.X.152.184:1521/ORCL
- Initial load database username: **admin**
- Initial load database password: <**password**>
- Check "Use different credentials for replication" and provide **ggadmin** and **password**.
- Don't check create private endpoint option.

Press **Create**

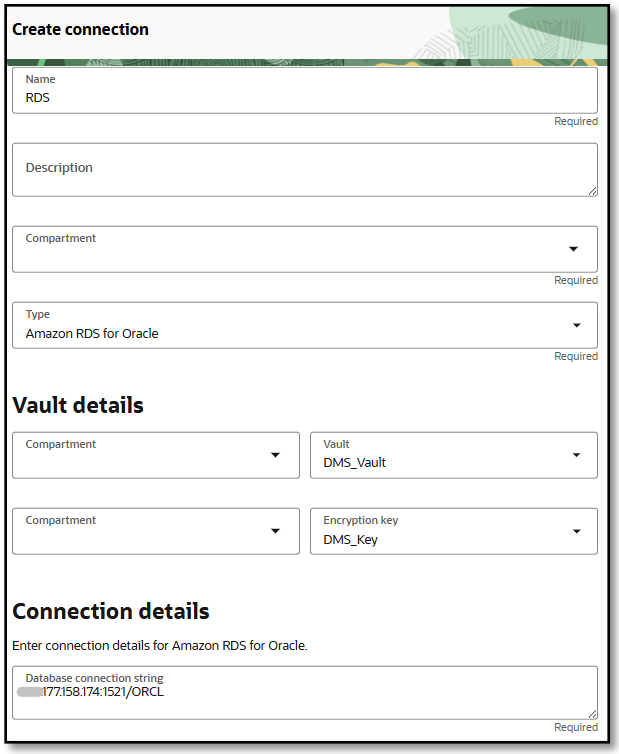

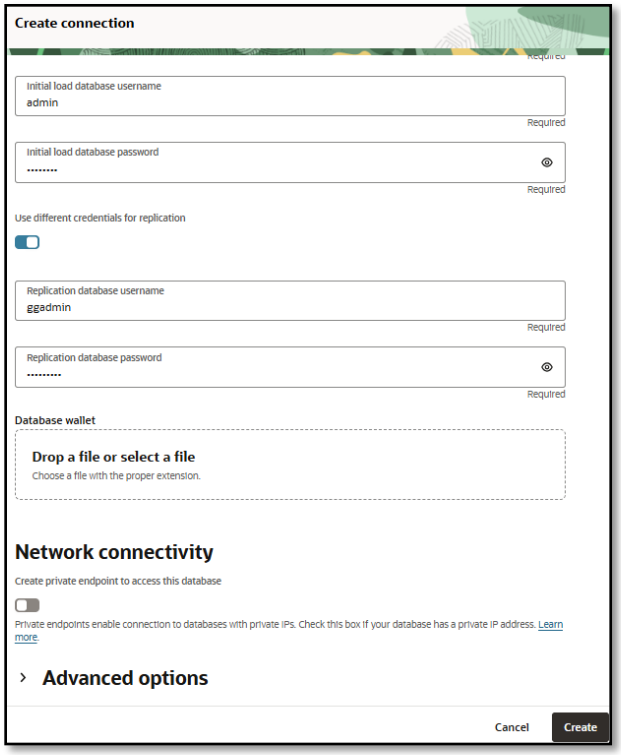

Once your newly created connection is in Active state, test it by clicking "Actions >Test connection" :

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

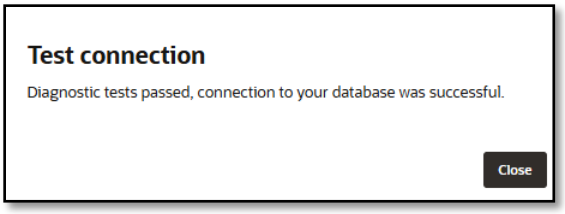

Task 13: Create database connection for the target database

In the OCI Console Menu, go to Migration & Disaster Recovery> Database Migration > Database Connections.

Press Create connection.

On the section Database Details, fill in the following entries, otherwise leave defaults:

- Name: **TargetATP**
- Type: **Oracle Autonomous Database**
- Vault: **DMS_Vault**
- Encryption Key: **DMS_Key**

  Select the Autonomous database name in your compartment i.e: **TargetATP**
- Initial load database username: **admin**
- Initial load database password: <**Admin password**>
- Check "Use different credentials for replication" and provide **ggadmin** and **password**.
- Network connectivity: **Create** private endpoint to access this database. Select the correct subnet.

Press **Create**

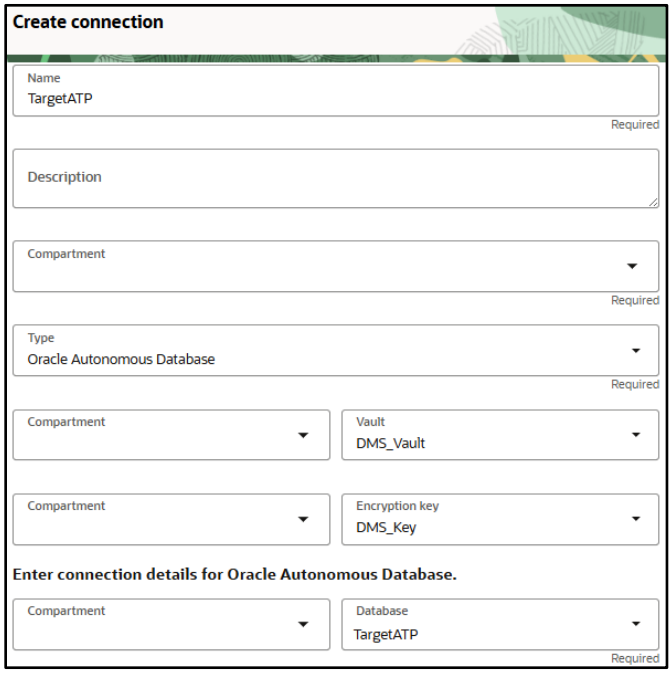

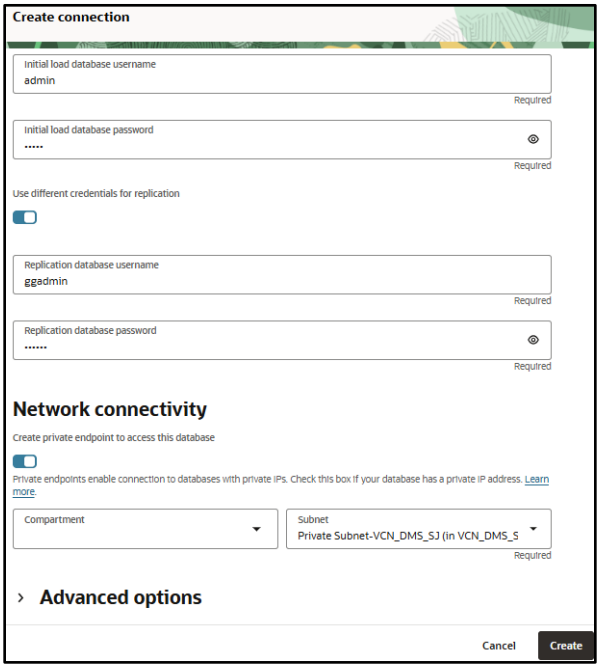

Test your connection as in the previous task.

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

Task 14: Create Migration

In the OCI Console Menu, go to Migration & Disaster Recovery> Database Migration > Migrations.

Press **Create Migration**.

On the page **Add Details**, fill in the following entries, otherwise leave defaults:

- Name: **TestMigration**
- Source database connection: **RDS**
- Target Database connection: **TargetATP**
- Transfer medium for initial load: **Datapump via Amazon Simple Storage Service**

Enter the Amazon S3 bucket information:

- Name:<your bucket name>
- Region:<bucket region>
- Key id:< AWS Access Key ID>
- Access key: < Secret Access Key>

- Export directory object name: **dumpdir**
- Object Storage Bucket: **DMSStorage**
- **Check** Use Online Replication

Click Create

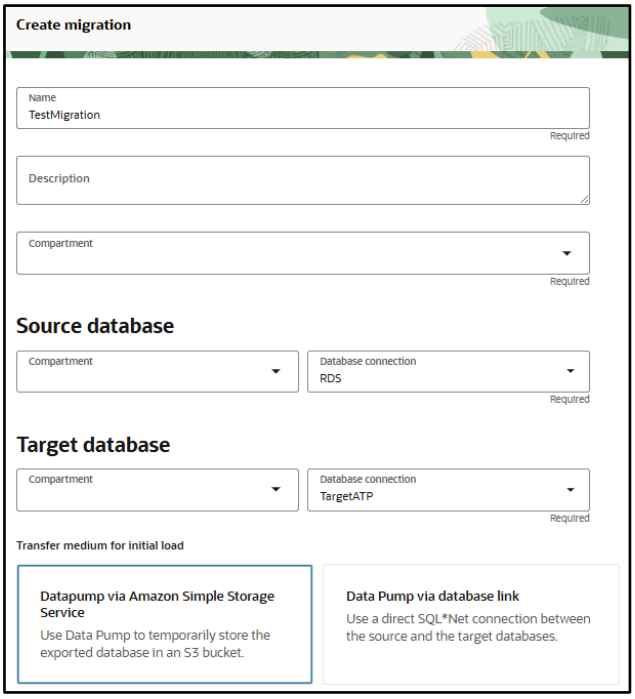

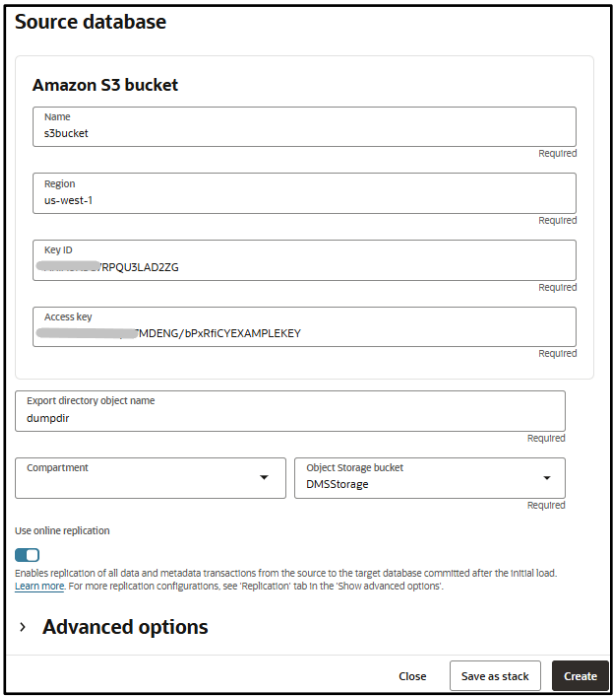

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

Task 15: Validate Migration

In this step you will validate a migration prior to running it. It will check the connections and settings for the source and target. Cloud premigration advisor (CPAT) will look for source and target incompatibilities.

In the OCI Console Menu, go to Migration & Disaster Recovery > Database Migration > Migrations.

Select **TestMigration**.

If the migration is still being created, wait until Lifecycle State is Active.

Press **Validate** button.

Click on the **Jobs** tab.

Click on most recent Evaluation Job

Click on the **Phases** tab.

Phases will be shown, and status will be updated as phases are completed. It can take 2 minutes before the first phase is shown.

If a phase has failed, it will show with status **Failed**. In this case press Actions>**Download Log** to learn more about the reason of failure.

Click **Run premigration advisor** phase name to open the Validation premigration advisor detail page (You should not find issues in this exercise but below lines would walk you thru an event when the phase fails). From this page you can download the CPAT report, view the report statistics, and drill down in the Checks list as shown:

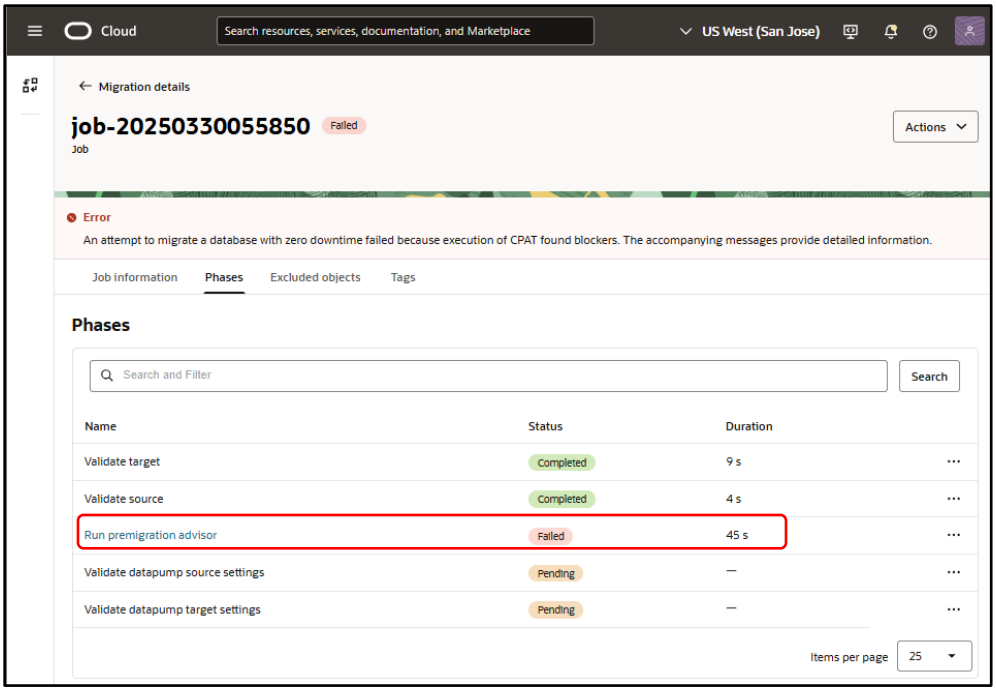

You can still download the advisor report as a text file, but now you can also navigate through the different checks. The summary view is displayed as follows:

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

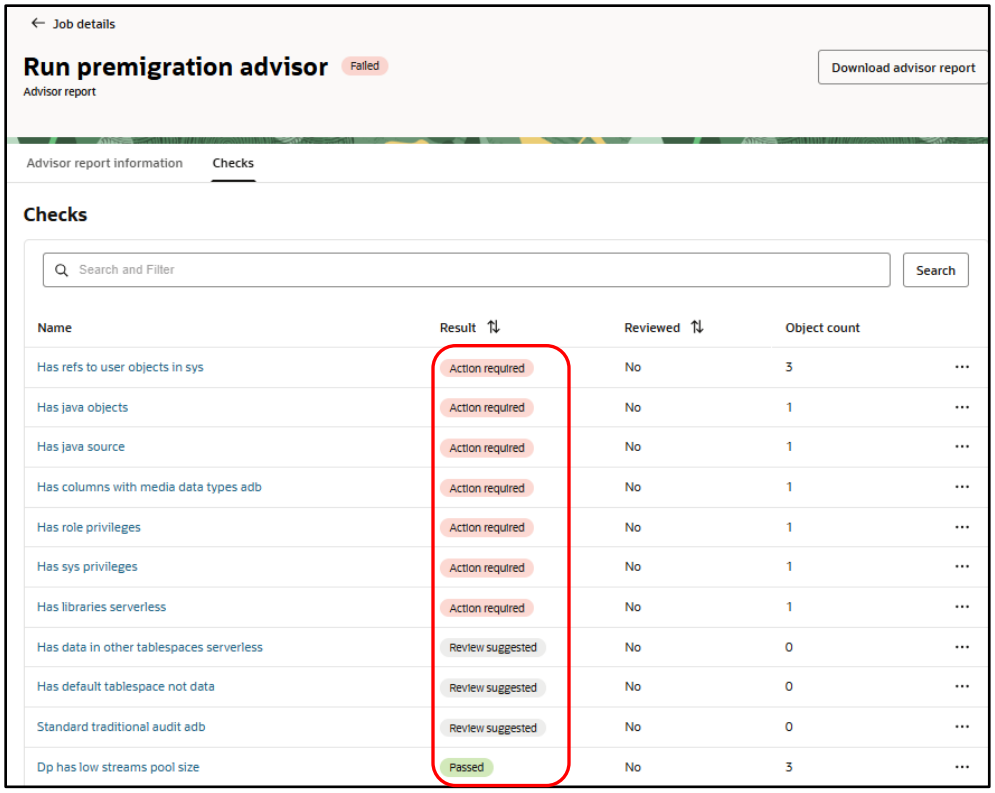

You can click a check name in the list to display details about a specific check from the CPAT report. You can mark a check as **Reviewed** or **Unreviewed**, this state is only for your convenience to track each check. For certain checks, CPAT generates a remedial script on the file system of the source database server. You can run the script on the source database to resolve the issue identified by the check. The checks page will also let you filter by this state (left side of screen):

The **View check details** panel is displayed as follows:

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

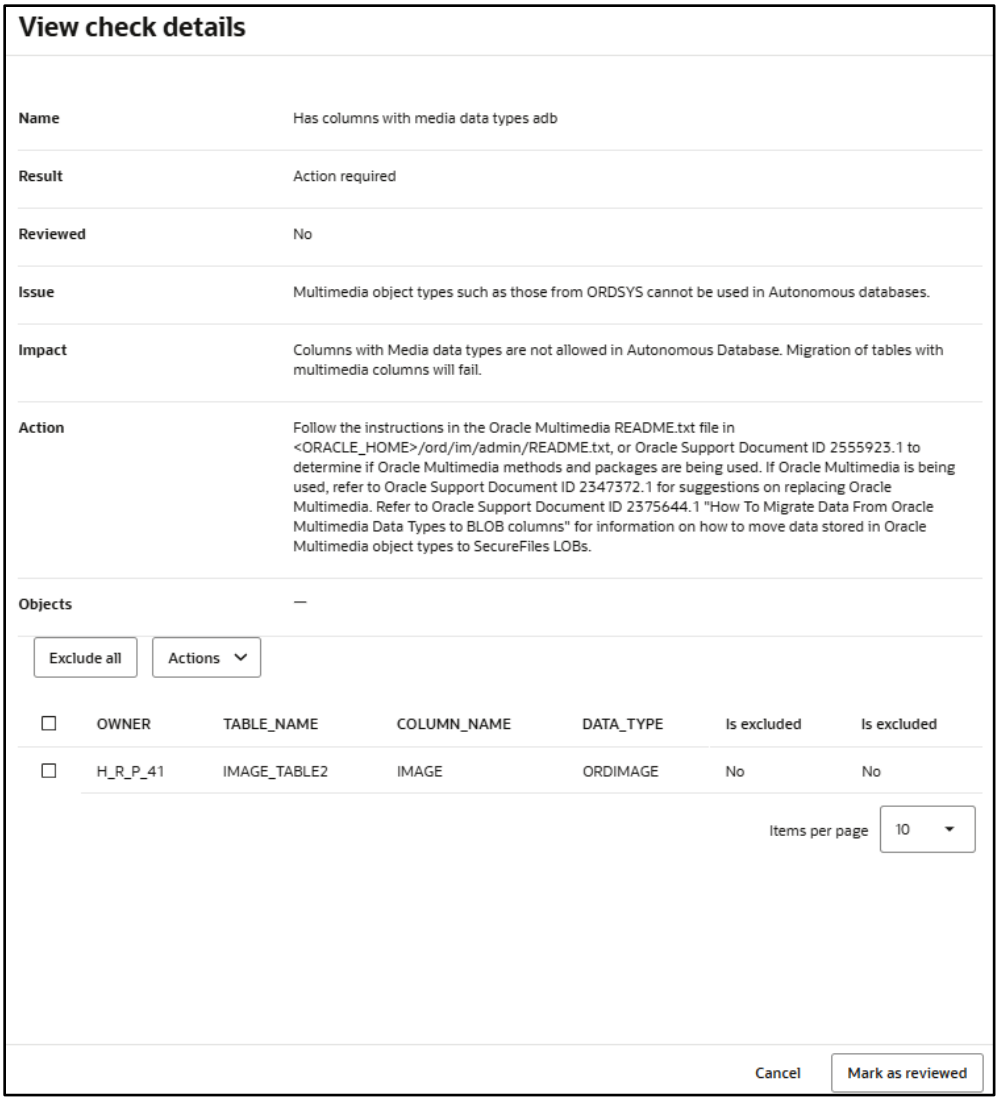

Once you have cleared all "Action Required" checks then the validation Job can be run again. Repeat the process until **Validate premigration advisor** phase completes with no error as shown:

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

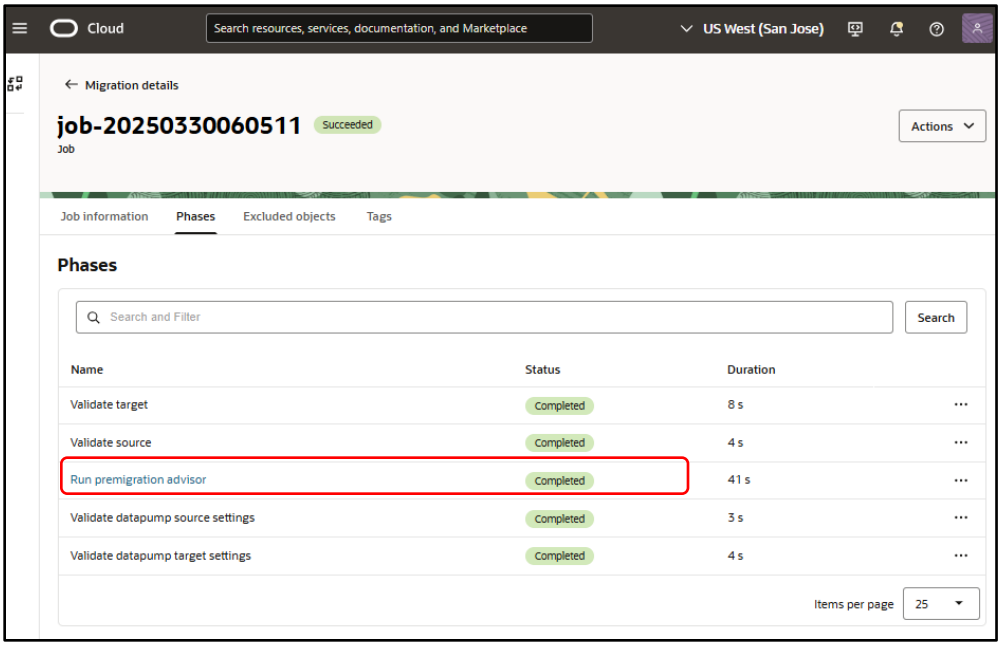

Task 16: Run Migration

After a successful validation, a Migration can be run to perform the data transfer.

In the OCI Console Menu, go to Migration & Disaster Recovery > Database Migration > Migrations.

Select **TestMigration**.

Press Start to begin the migration.

A confirmation dialog opens, and the job can be configured to pause at any point by selecting a phase in Require User Input After, the pre-selected value is Monitor replication lag. This phase monitors Oracle GoldenGate Extract and Replicat operations until Replicat has caught up on the target database; end-to-end (E2E) replication lag should be less than 30 seconds.

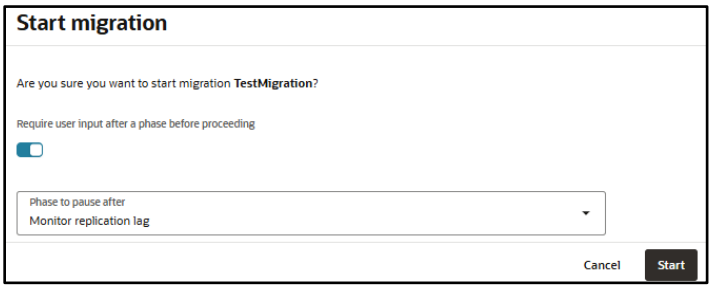

When the selected phase completes, the job will enter in a Waiting state until it is resumed (or terminated). If it was selected to pause after the phase Monitor Replication Lag, the transaction replication continues during the Waiting state. It will stop upon resume.

This is the point where a migration user would stop the source application so that no more transactions are applied to the source DB. You can now press **Resume** on the job to complete replication.

In the Resume Job dialog, chose the **Switchover** phase and press **Resume**. The Switchover phase will gracefully stop replication and signal the target application to initiate transactions to the target DB. Find more information about the switchover phase in our documentation.

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

After the phase Switchover has completed, the workload on the target database (end of downtime) can start.

The last phase is Cleanup, click on the Resume button and click again Resume on the phase selection window. This phase performs cleanup operations such as deleting GoldenGate Extract and GoldenGate Replicat processes and connection details on source and target database respectively, removing Autonomous Database access to wallet, and so on. Learn more of the different phases at the following link.

The migration runs the final cleanup phases and shows as Succeeded when finished.

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

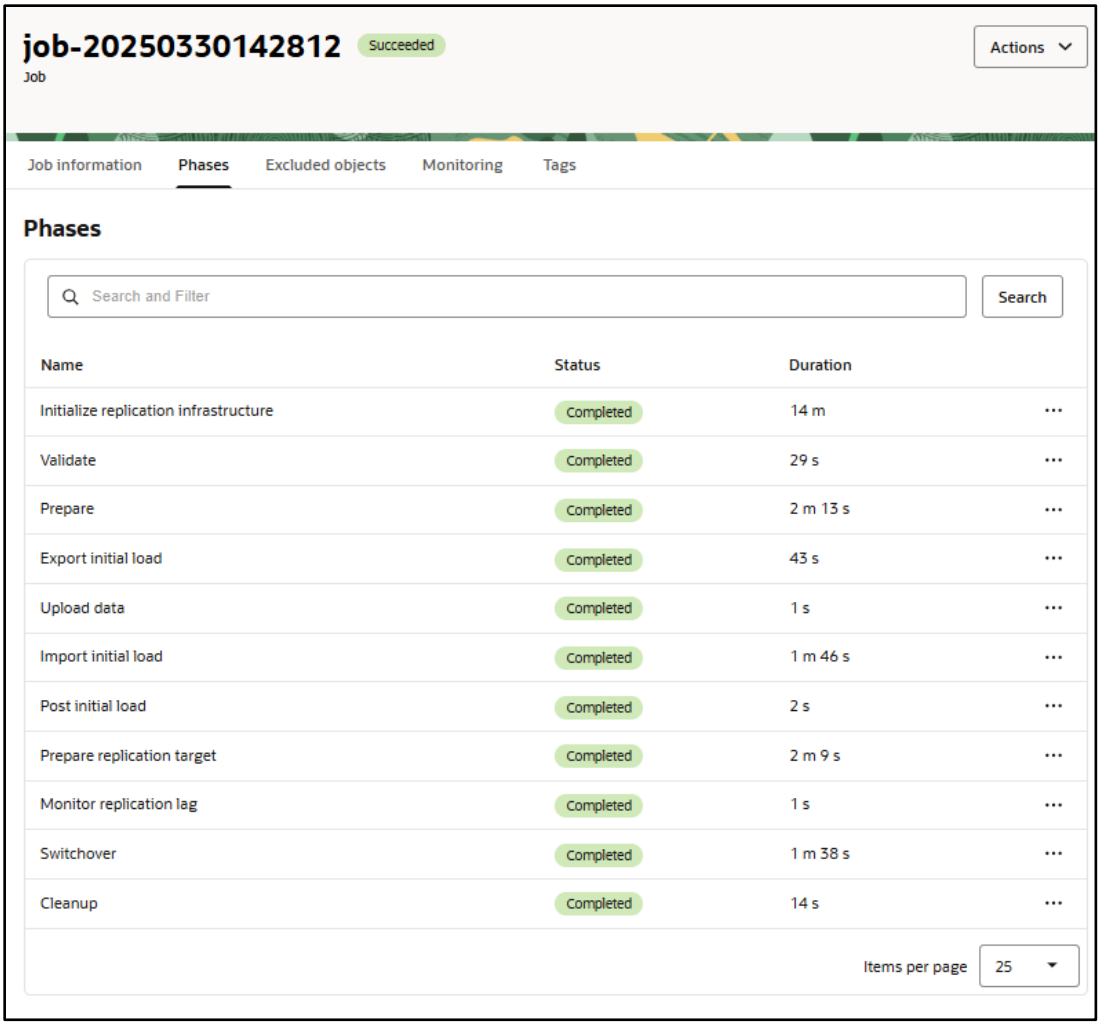

Your migration is now completed.!

Issues faced while performing this guide

During the Migration phase on DMS I got the following error:

"PRGD-1042 : query to retrieve information from database view rdsadmin.rds_file_util.listdir failed

PRGD-1002 : SELECT statement "SELECT FILENAME FROM TABLE(rdsadmin.rds_file_util.listdir(p_directory => 'DATA_PUMP_DIR')) WHERE FILENAME LIKE '%ZDM_502_DP_EXPORT_9642_dmp_%'" execution as user "admin" failed for database with Java Database Connectivity (JDBC) URL "jdbc:oracle:thin:@(description=(address=(protocol=tcp)(port=1521)(host=54.177.158.174))(connect_data=(service_name=ORCL)))"

ORA-20199: Error in rdsadmin.rds_file_util."

Copyright © 2025, Oracle and/or its affiliates / Dropdown Options

This issue was identified in the **AWS side**, and I got help from a support engineer: “there is an issue identified which is causing “RDSADMIN.RDS_FILE_UTIL.LISTDIR” package failure. Run the workaround command “exec rdsadmin.rdsadmin_rman_util.validate_tablespace('USERS'); “ you will be able to list the files successfully with RDS_FILE_UTIL.LISTDIR.”

This allowed the migration to progress.

Connect with us

Call +**1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

blogs.oracle.com facebook.com/oracle twitter.com/oracle

Copyright © 2025, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120

Disclaimer: If you are unsure whether your data sheet needs a disclaimer, read the revenue recognition policy. If you have further questions about your content and the disclaimer requirements, e-mail REVREC_US@oracle.com.