Oracle® Database

Step by Step Data Replication Using Oracle GoldenGate Microservices Architecture

21c

F41988-04

August 2022

Oracle Database Step by Step Data Replication Using Oracle GoldenGate Microservices Architecture, 21c

F41988-04

Primary Author: Oracle Corporation

Contents

Preface

1 Quickstarts

2 Configure Deployments

3 Manage Deployments from the Service Manager

4 Configure Data Replication Processes from the Administration Service

5 Configure Paths to Transport Trail Data

Preface

The *Step by Step Data Replication Using Oracle GoldenGate Microservices Architecture* is a walk through of the entire Oracle GoldenGate data replication cycle using Microservices.

- Audience
- Documentation Accessibility
- Conventions
- Related Information

Audience

This guide is intended for administrators and users who are familiar with Oracle GoldenGate concepts and architecture and who are interested in learning to use the microservices for performing various Oracle GoldenGate data replication tasks.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc`.

Accessible Access to Oracle Support

Oracle customers who have purchased support have access to electronic support through My Oracle Support. For information, visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info` or visit `http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs` if you are hearing impaired.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, such as "From the File menu, select **Save**." Boldface also is used for terms defined in text or in the glossary.
italic *`italic`*	Italic type indicates placeholder variables for which you supply particular values, such as in the parameter statement: `TABLE` *`table_name`*. Italic type also is used for book titles and emphasis.

Convention	Meaning
`monospace` `MONOSPACE`	Monospace type indicates code components such as user exits and scripts; the names of files and database objects; URL paths; and input and output text that appears on the screen. Uppercase monospace type is generally used to represent the names of Oracle GoldenGate parameters, commands, and user-configurable functions, as well as SQL commands and keywords.
UPPERCASE	Uppercase in the regular text font indicates the name of a process or utility unless the name is intended to be a specific case. Keywords in upper case (`ADD EXTRACT`, `ADD EXTTRAIL`, `FORMAT RELEASE`).
LOWERCASE	Names of processes to be written in lower case. Examples: `ADD EXTRACT exte`, `ADD EXTRAIL ea`.
{ }	Braces within syntax enclose a set of options that are separated by pipe symbols, one of which must be selected, for example: {*option1* | *option2* | *option3*}.
[ ]	Brackets within syntax indicate an optional element. For example in this syntax, the `SAVE` clause is optional: `CLEANUP REPLICAT` *group_name* `[, SAVE` *count*`]`. Multiple options within an optional element are separated by a pipe symbol, for example: [*option1* | *option2*].
Sample Locations	Compass directions such as east, west, north, south to be used for demonstrating Extract and Replicat locations. Datacenters names to use the standard similar to `dc1`, `dc2`.
Group names	Prefixes for each process, as follows: • Extract: ext. Usage with location: `extn`, where ***n*** indicates 'north' compass direction. • Replicat: rep. Usage with location: `repn`, where ***n*** indicates 'north' compass direction. • Distribution Path: `dp`. Usage with location: `dpn`, where ***n*** indicates 'north' compass direction. • Checkpoint table: `ggs_checkpointtable` • Trail file names: e or d depending on whether the trail file is for the Extract of distribution path. Suffix derived in alphabetical order. Usage for an Extract trail file: `ea`, `eb`, `ec`. • Trail file subdirectory: The name will use compass directions to refer to the trail subdirectories. Example for trail subdirectory name would be `/east`, `/west`, `/north`, `/south`.

Related Information

The Oracle GoldenGate Product Documentation Libraries are found at:

Oracle GoldenGate Documentation

Oracle GoldenGate for Big Data Documentation:

https://docs.oracle.com/en/middleware/goldengate/big-data/index.html

For additional information on Oracle GoldenGate, refer to:

https://www.oracle.com/middleware/technologies/goldengate.html

Oracle Database High Availability

1 Quickstarts

1.1 Set Up Data Replication with Oracle GoldenGate Microservices Architecture

Use this quickstart to configure data replication using Oracle GoldenGate Microservices Architecture for a multitenant container database with two pluggable databases to demonstrate data replication from an Oracle to Oracle database in a HUB configuration.

> **Note:**
>
> This quickstart does not perform an initial load instantiation and assumes that the tables and data are the same in the source and target endpoints.

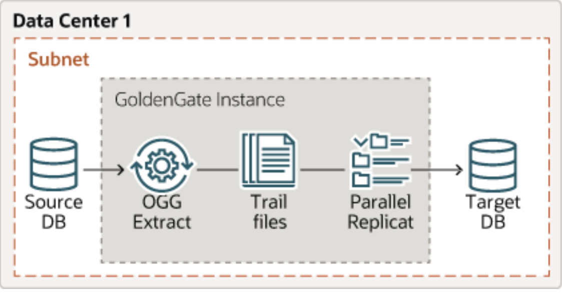

The source and target databases in this diagram refer to the container and pluggable databases (PDBs).

Container Database (CDB$ROOT) Process Names	Pluggable Database (DBEAST) Process Names	Pluggable Database (DBWEST) Process Name
• **`CDB$ROOT` database user: `c##ggadmin`** • **Database credential alias: `cggnorth`**	• Database user: **`ggadmin`** • Database alias: **`ggeast`** • Extract: **`exte`**	• Database user: **`ggadmin`** • Database alias: **`ggwest`** • Extract: **`extw`**

Configure and Set Privileges for Oracle Multitenant Database

In Oracle database, you need to enable replication for Oracle GoldenGate and assign privileges to the database user at the CDB level and the pluggable database (PDB) level.

The database is in `ARCHIVELOG` mode and `FORCE LOGGING` and Supplemental Logging is enabled. For the container database, assign the following privileges to the common user (`cdb$root`):

CDB User Privileges

```
## CGGNORTH DATABASE SETUP AT CDB LEVEL
alter session set container=cdb$root;
alter system set enable_goldengate_replication=TRUE;
alter system set streams_pool_size=2G;
alter database force logging;
alter database add supplemental log data;
archive log list;
create tablespace GG_DATA datafile '+DATA' size 100m autoextend on
next 100m;
create user c##ggadmin identified by Password container=all default
tablespace GG_DATA temporary tablespace temp;
grant alter system to c##ggadmin container=all;
grant dba to c##ggadmin container=all;
grant create session to c##ggadmin container=all;
grant alter any table to c##ggadmin container=all;
grant resource to c##ggadmin container=all;
exec
dbms_goldengate_auth.grant_admin_privilege('c##ggadmin',container=>'all
');
```

Source PDB User Privileges (`DBEAST`)

```
alter session set container=DBEAST;
create tablespace GG_DATA datafile '+DATA' size 100m autoextend on
next 100m;
create user ggadmin identified by Password container=current;
grant create session to ggadmin container=current;
grant alter any table to ggadmin container=current;
grant resource to ggadmin container=current;
grant dba to ggadmin container=current;
exec dbms_goldengate_auth.grant_admin_privilege('ggadmin');
```

Target PDB User Privileges (`DBWEST`):

```
alter session set container=DBWEST;
create user ggadmin identified by password container=current;
grant alter system to ggadmin container=current;
grant create session to ggadmin container=current;
grant alter any table to ggadmin container=current;
grant resource to ggadmin container=current;
grant dba to ggadmin container=current;
grant dv_goldengate_admin, dv_goldengate_redo_access to ggadmin
container=current;
exec dbms_goldengate_auth.grant_admin_privilege('ggadmin');
```

Note:

Granting DBA role is not mandatory for every user. Privileges should be granted depending on the actions that the user needs to perform on the database. For example, to grant DML operation privileges to insert, update, and delete transactions to `ggadmin`, use the `GRANT ANY INSERT/UPDATE/DELETE` privileges and to further allow users to work with tables and indexes as part of DML operations, use the `GRANT CREATE/DROP/ALTER ANY TABLE/INDEX` privileges. In this quickstart, the assumption is that the database user is a database administrator. See Oracle Database Privileges and Other Requirements for Multitenant Container Databases to know more about specific privilege requirements.

Configure the Replication Process from Oracle GoldenGate MA Web Interface

Data replication processes include Extracts, Replicats processes, along with Distribution Paths (`DISTPATH`) or Receiver Paths or Target-Initiated Paths (`RECVPATH`).

Using the following steps, you'll be able to configure data capture (Extract) and apply (Replicat) processes. You'll also be able to test if the replication has started. The `DISTPATH` process is not used for this configuration.

Step 1: Add Database Credentials from the Administration Service

In this section, you'll add the database credentials to connect to the source and target databases using EZConnect.

1. Keep your database user credentials, which created in the previous session, ready. You'll use them to connect Oracle GoldenGate to the database server.
2. Open the Service Manager login page in a web browser and log in to the Service Manager with your Oracle GoldenGate administrator user credentials. If logging in for the first time, you have to log in with the administrator account user credentials, created when adding your deployment with Oracle GoldenGate Configuration Assistant wizard.
3. From the Service Manager Overview page, click the port number for the Administration Service of the deployment.

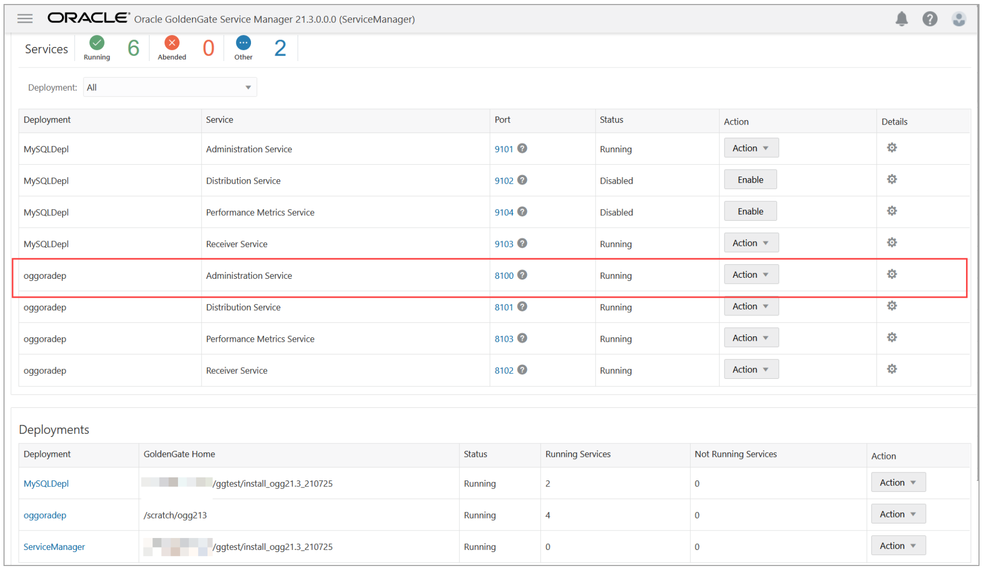
ORACLE Oracle GoldenGate Service Manager 21.3.0.0.0 (ServiceManager)
Services
Running 6
Abended 0
Other 2
Deployment: All

Deployment	Service	Port	Status	Action	Details
MySQLDepl	Administration Service	9101	Running	Action	⚙
MySQLDepl	Distribution Service	9102	Disabled	Enable	⚙
MySQLDepl	Performance Metrics Service	9104	Disabled	Enable	⚙
MySQLDepl	Receiver Service	9103	Running	Action	⚙
oggoradep	Administration Service	8100	Running	Action	⚙
oggoradep	Distribution Service	8101	Running	Action	⚙
oggoradep	Performance Metrics Service	8103	Running	Action	⚙
oggoradep	Receiver Service	8102	Running	Action	⚙

Deployments

Deployment	GoldenGate Home	Status	Running Services	Not Running Services	Action
MySQLDepl	[redacted]/ggtest/install_ogg21.3_210725	Running	2	0	Action
oggoradep	/scratch/ogg213	Running	4	0	Action
ServiceManager	[redacted]/ggtest/install_ogg21.3_210725	Running	0	0	Action

This opens the Administration Service login page.

4. Log in to the Administration Service using the same credentials, which you used to log in to the Service Manager. The Administration Service **Overview** page is displayed.

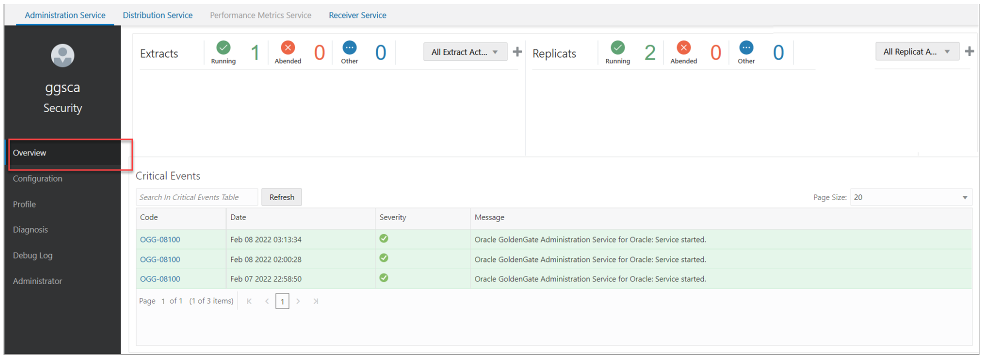

Code	Date	Severity	Message
OGG-08100	Feb 08 2022 03:13:34	✓	Oracle GoldenGate Administration Service for Oracle: Service started.
OGG-08100	Feb 08 2022 02:00:28	✓	Oracle GoldenGate Administration Service for Oracle: Service started.
OGG-08100	Feb 07 2022 22:58:50	✓	Oracle GoldenGate Administration Service for Oracle: Service started.

Page 1 of 1 (1 of 3 items)

5. Click the **Application Navigation** icon to open the left-navigation pane and click **Configuration** to open the **Database** tab of the **Configuration** page.

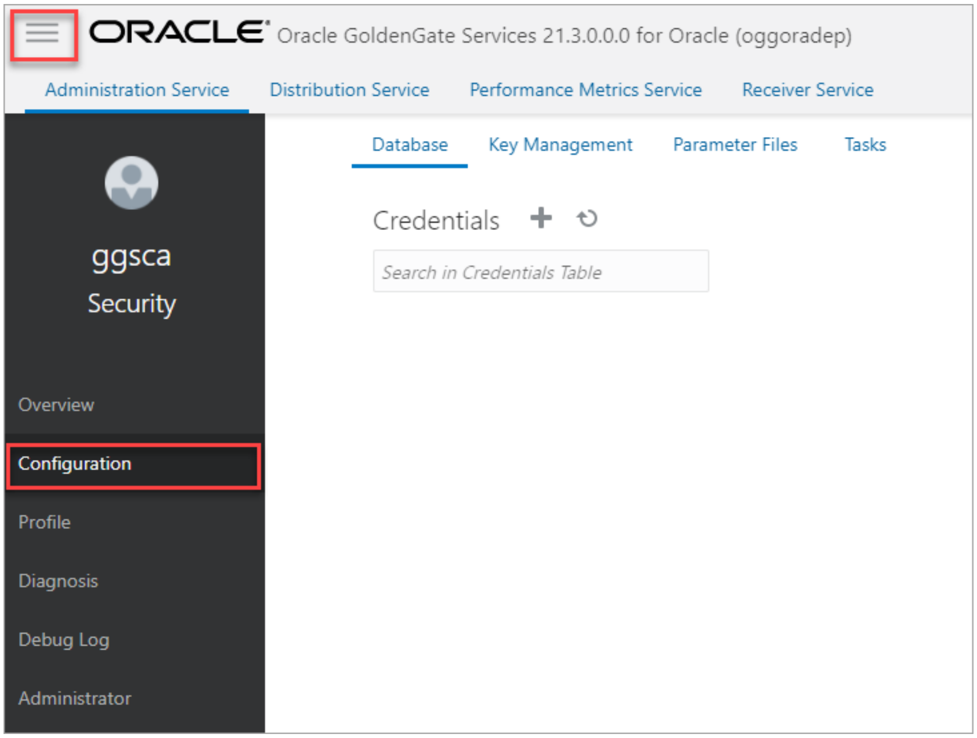

6. Click the plus (+) sign in the **Credentials** section to begin adding database user credentials.

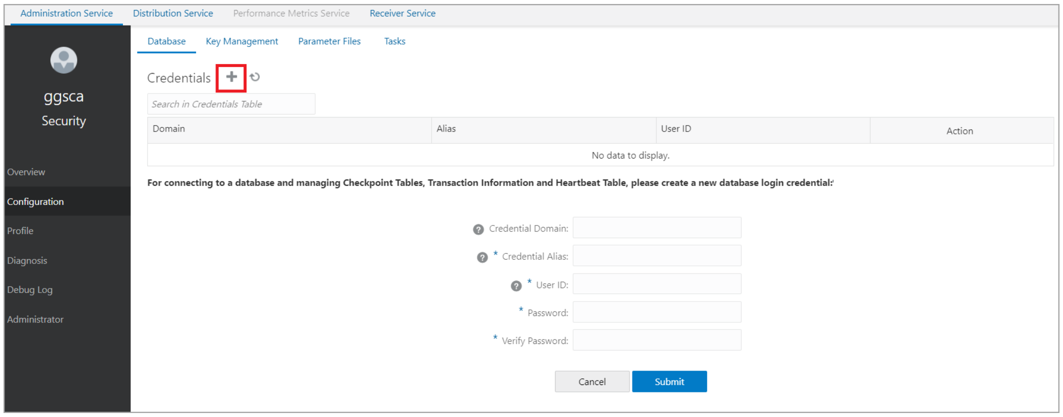

Domain	Alias	User ID	Action
No data to display.			

For connecting to a database and managing Checkpoint Tables, Transaction Information and Heartbeat Table, please create a new database login credential:

Credential Domain:

* Credential Alias:

* User ID:

* Password:

* Verify Password:

7. You need to add connections for container database (CDB) and pluggable databases (PDBs). Each CDB is used to capture (Extract) from the source database and PDB for delivery (Replicat).

   Use the EZconnect syntax to configure the database connection. You need the username, password, hostname, port number, and service name connection information to use the EZConnect syntax.

   Here's the syntax that you need to specify in the User ID field:

   username@*hostname*:*port*/*service_name*

   Here's an example for setting the User ID with EZConnect:

   ```
   c##ggadmin@dc.example.com:1521/dc1.example.com
   ```

The following screen shows the database credential (**cggnorth**) for connecting to the user `c#ggadmin` added to the credentials list. You can also see the credentials being added for the **Alias** (`ggeast`) for connecting to the `DBEAST` PDB.

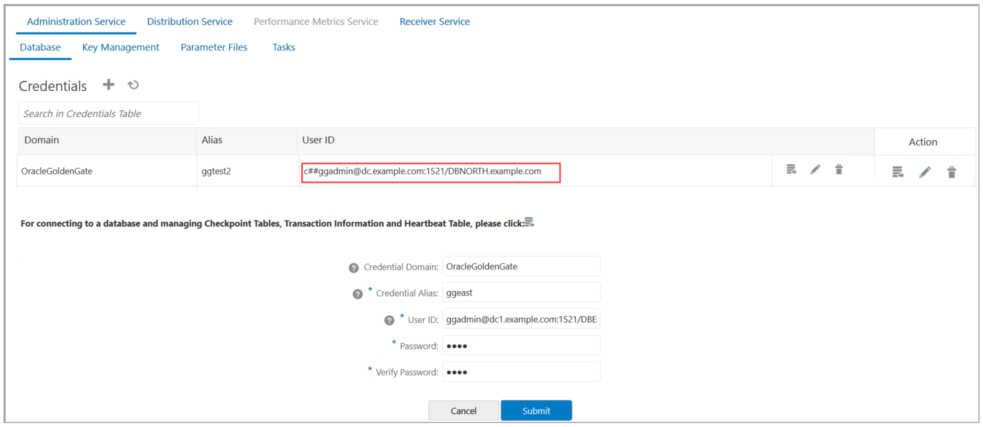

Domain	Alias	User ID	Action
OracleGoldenGate	ggtest2	c##ggadmin@dc.example.com:1521/DBNORTH.example.com	

For connecting to a database and managing Checkpoint Tables, Transaction Information and Heartbeat Table, please click:

Credential Domain:	OracleGoldenGate
* Credential Alias:	ggeast
* User ID:	ggadmin@dc1.example.com:1521/DBE
* Password:	••••
* Verify Password:	••••

8. Click on the blue icon in the **Actions** column to connect to the database. The icon turns blue when the connection is successful.

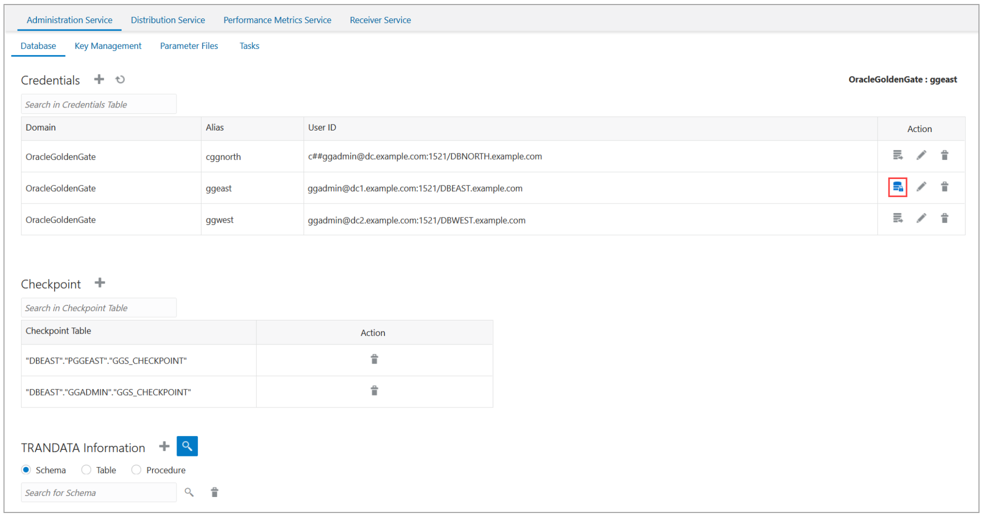

Domain	Alias	User ID	Action
OracleGoldenGate	cggnorth	c##ggadmin@dc.example.com:1521/DBNORTH.example.com	
OracleGoldenGate	ggeast	ggadmin@dc1.example.com:1521/DBEAST.example.com	
OracleGoldenGate	ggwest	ggadmin@dc2.example.com:1521/DBWEST.example.com	

Checkpoint Table	Action
"DBEAST"."PGGEAST"."GGS_CHECKPOINT"	
"DBEAST"."GGADMIN"."GGS_CHECKPOINT"	

After connecting to the database, the sections to add checkpoint table, TRANDATA, and heartbeat table are displayed.

Step 2: Add TRANDATA, Heartbeat, and Checkpoint Tables

In this section, you will add TRANDATA for the source database to enable writing information to the redo logs. This would ensure that the rows added to the source database are uniquely identified on the target database and are updated. You'll add heartbeat tables for the source and target databases to monitor any possible lags. You will also add a checkpoint table for the target database to ensure that if there is a failure, then the Extract and Replicat processes can restart from the point of failure.

1. Add TRANDATA to the source connection. Use the **TRANDATA Information** section to set up database logging properties. This is an essential step to enable supplemental logging and ensure that the data is written to the database redo log.

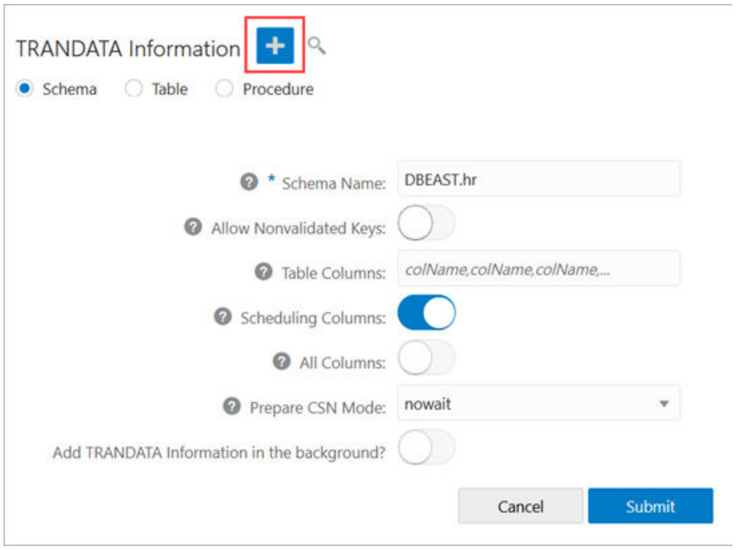

TRANDATA Information

● Schema ○ Table ○ Procedure

* Schema Name:	DBEAST.hr
Allow Nonvalidated Keys:	Off
Table Columns:	colName,colName,colName,...
Scheduling Columns:	On
All Columns:	Off
Prepare CSN Mode:	nowait
Add TRANDATA Information in the background?	Off

After you add the trandata, you can search for the schema for which you've add the trandata, using the search icon. This will display the trandata information. The following image shows the trandata information for the `HR` schema in the pluggable database `DBEAST`.

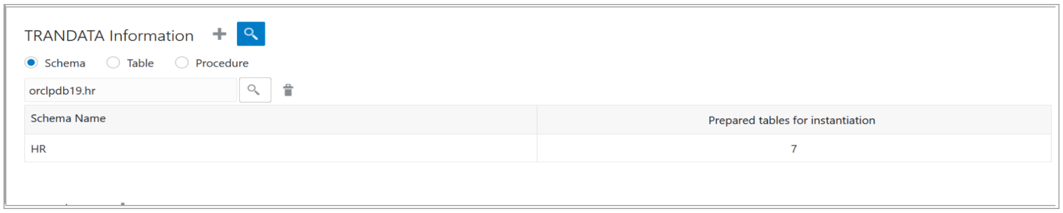

TRANDATA Information

● Schema ○ Table ○ Procedure

orclpdb19.hr

Schema Name	Prepared tables for instantiation
HR	7

See Configuring Logging Properties to learn the steps for configuring the logging properties at the Schema, Table or Procedure level.

2. To set up the checkpoint table for Replicat, you need to connect to the target database credentials (**ggwest**) from the Credentials section.

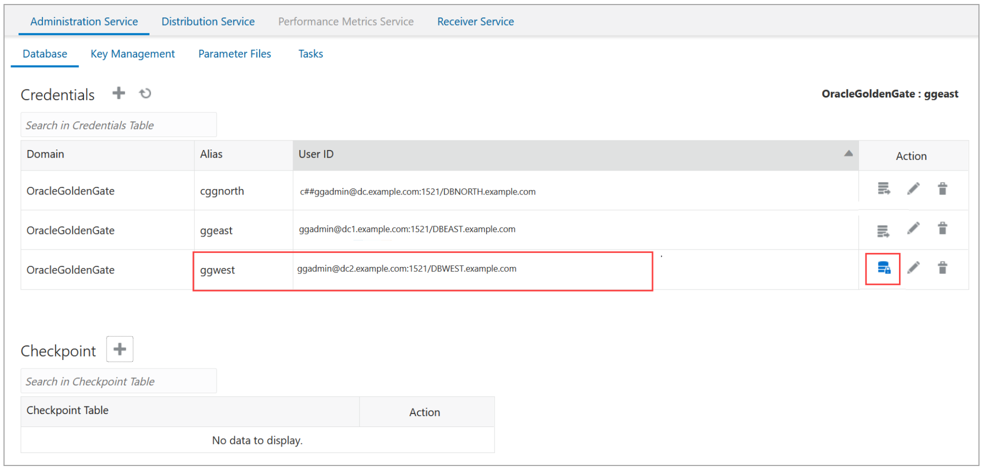

Credentials

OracleGoldenGate : ggeast

Search in Credentials Table

Domain	Alias	User ID	Action
OracleGoldenGate	cggnorth	c##ggadmin@dc.example.com:1521/DBNORTH.example.com	
OracleGoldenGate	ggeast	ggadmin@dc1.example.com:1521/DBEAST.example.com	
OracleGoldenGate	ggwest	ggadmin@dc2.example.com:1521/DBWEST.example.com	

Checkpoint

Search in Checkpoint Table

Checkpoint Table	Action
No data to display.	

3. Click the plus sign (+) to add the checkpoint table for the target (pluggable) database.

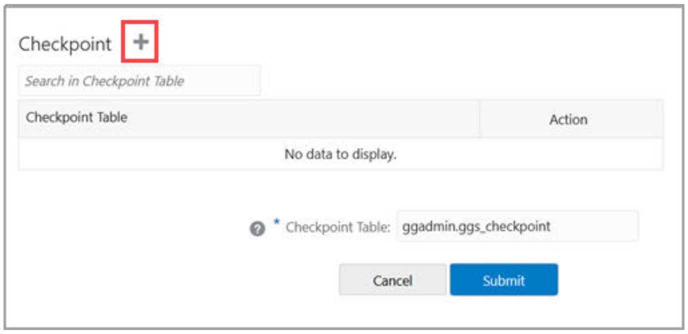

Checkpoint

Search in Checkpoint Table

Checkpoint Table	Action
No data to display.	

* Checkpoint Table: ggadmin.ggs_checkpoint

Click **Submit**. The checkpoint table is added.

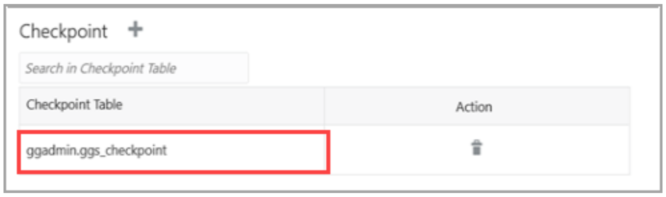

Checkpoint

Search in Checkpoint Table

Checkpoint Table	Action
ggadmin.ggs_checkpoint	

Also see the Before Creating an Extract section, for details on creating heartbeat tables.

4. Add the heartbeat tables for both source and target endpoints by connecting to **ggeast** and **ggwest** database credential aliases. Add the heartbeat table by clicking the plus sign.

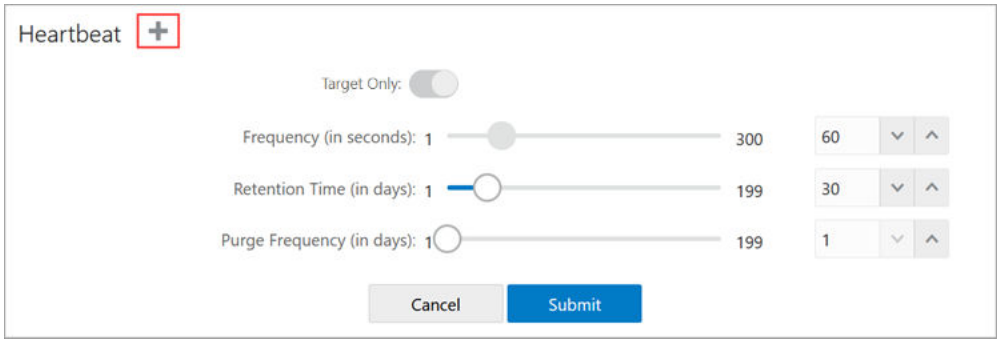

Target Only: ○

Frequency (in seconds):	1	300	60	⌄ ⌃
Retention Time (in days):	1	199	30	⌄ ⌃
Purge Frequency (in days):	1	199	1	⌄ ⌃

5. Click **Submit** after adjusting the heartbeat options.

Step 3: Add an Extract

In this section, you will add an Extract process (**exte**). The Extract process captures data from the source database and writes it to a trail file (**ea**).

1. Click the **Overview** option from the left-navigation pane of the Administration Service and click the plus sign (+) from the Extract section.

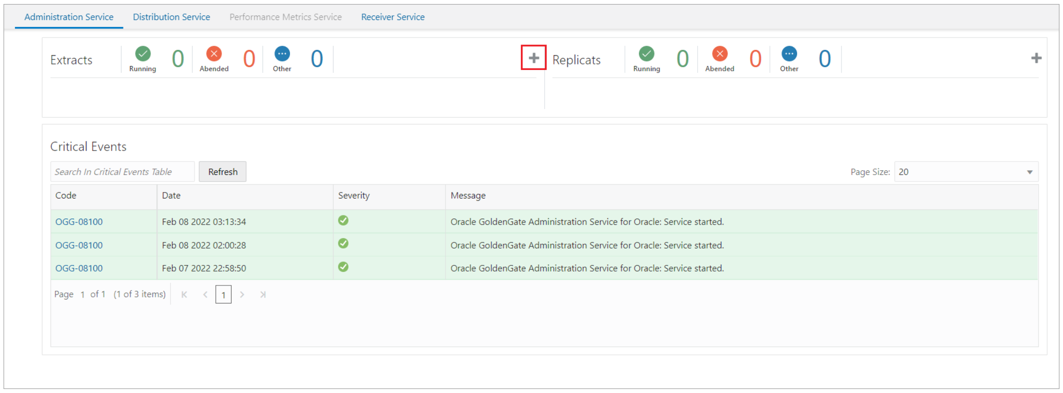

Code	Date	Severity	Message
OGG-08100	Feb 08 2022 03:13:34	✓	Oracle GoldenGate Administration Service for Oracle: Service started.
OGG-08100	Feb 08 2022 02:00:28	✓	Oracle GoldenGate Administration Service for Oracle: Service started.
OGG-08100	Feb 07 2022 22:58:50	✓	Oracle GoldenGate Administration Service for Oracle: Service started.

Page 1 of 1 (1 of 3 items)

2. From Add Extract wizard, select **Integrated Extract**.

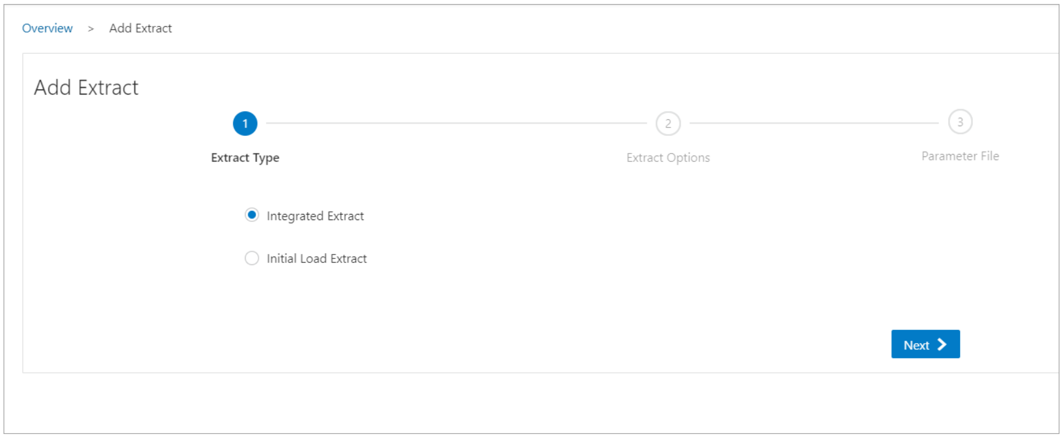

● Integrated Extract

○ Initial Load Extract

> **Note:**
>
> Before creating Replicat, you need to create an initial load Extract when starting the replication process for the first time. To learn about the initial load Extract and it's use case, see Instantiating Oracle GoldenGate with an Initial Load.

3. Click **Next** and specify the Extract options in the Extract Options screen. See the detailed steps to add an Extract from the Add an Extract section.

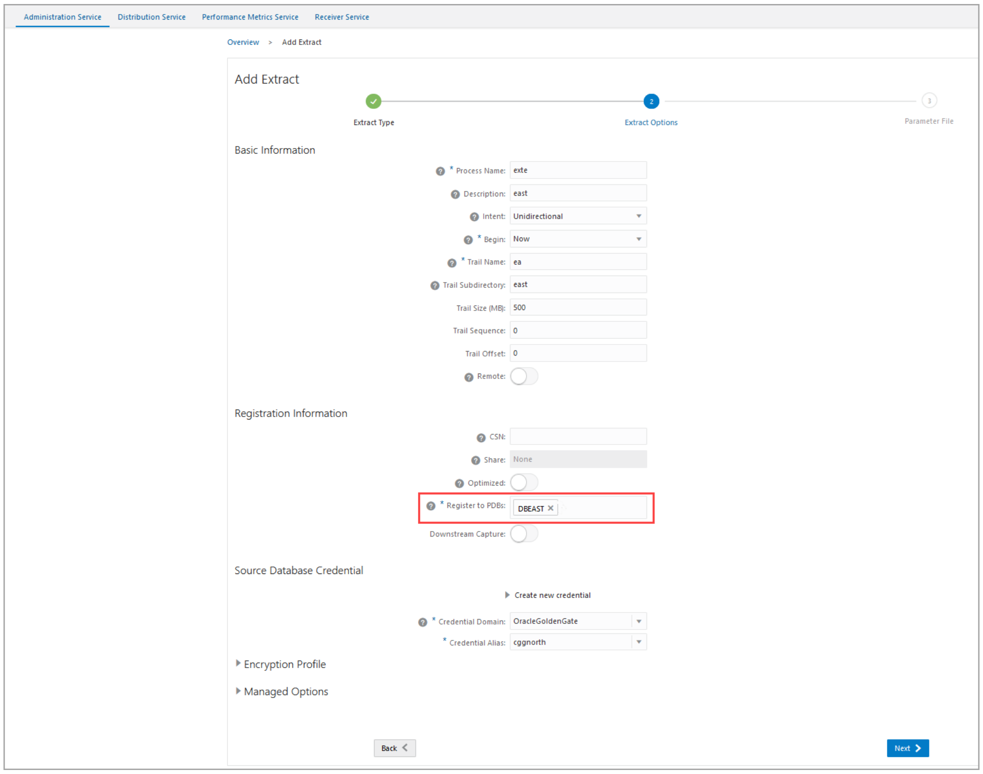

Add Extract

Basic Information

* Process Name: exte
Description: east
Intent: Unidirectional
* Begin: Now
* Trail Name: ea
Trail Subdirectory: east
Trail Size (MB): 500
Trail Sequence: 0
Trail Offset: 0
Remote: ○

Registration Information

CSN:
Share: None
Optimized: ○
* Register to PDBs: DBEAST ✕
Downstream Capture: ○

Source Database Credential

* Credential Domain: OracleGoldenGate
* Credential Alias: cggnorth

If you are creating the Extract for a pluggable database, then you'll see option **Register to PDBs** as soon as you enter the credentials domain and alias. Select the PDB in the container database that you want to use for replication.

4. After you enter the options for the Extract, click **Next**. The next screen displays the Extract parameter file to help you review the Extract settings.

   Here's the Extract parameter file for the Extract exte:

```
EXTRACT exte
USERIDALIAS cggnorth DOMAIN OracleGoldenGate
EXTTRAIL east/ea
SOURCECATALOG DBEAST
DDL INCLUDE MAPPED
TABLE hr.*;
```

   Review these settings and update the Extract configuration as needed.

For multitenant databases, you need to add entries for Extract to capture from multiple pluggable databases to a single trail. In the parameter file, source objects must be specified in `TABLE` and `SEQUENCE` statements with their fully qualified three-part names in the format of *`container.schema.object`* or using the `SOURCECATALOG` parameter with two-part names *`schema.object`*.

Click **Create and Run** to start your Extract.

Step 4: Add a Replicat

In this section, you will add a Replicat process (**repe**). The Replicat process delivers the change data from the trail file (**ea**) created by the Extract, to the target database. Replicat reads the trail file on the target database, reconstructs the DML or DDL operations, and applies them to the target database.

1. Before Creating a Replicat, make sure that you added your checkpoint table for the target database (DBWEST) by connecting to the **ggwest** database credentials.
2. Select a Replicat type to deliver data to the target database. Follow the wizard to complete adding a Replicat. See How to Add a Replicat.

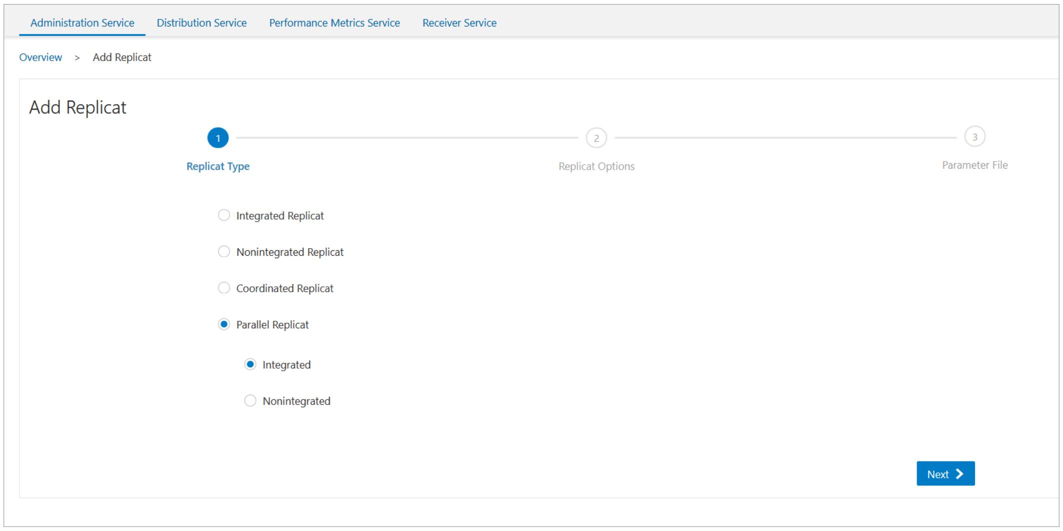

○ Integrated Replicat
○ Nonintegrated Replicat
○ Coordinated Replicat
● Parallel Replicat
    ● Integrated
    ○ Nonintegrated

3. Enter the **Parallel Nonintegrated Replicat** options in the Replicat Options screen.

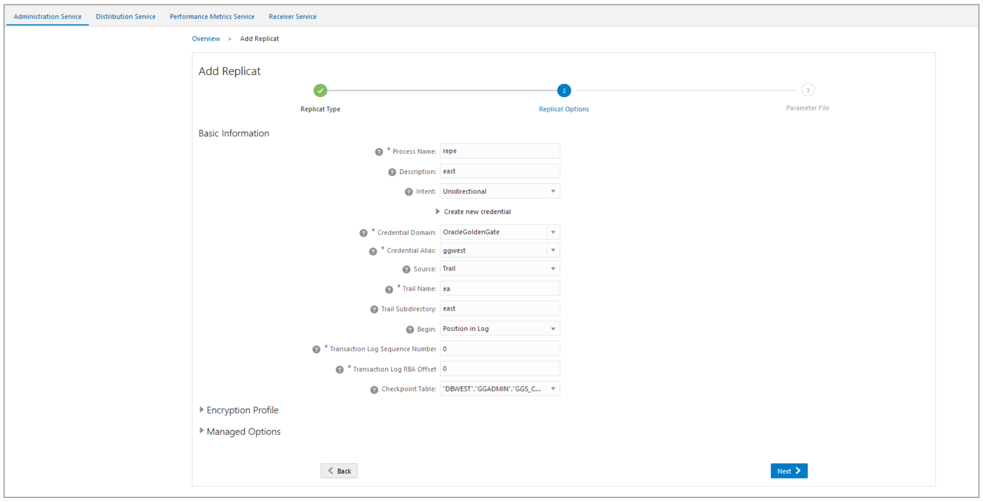

Add Replicat

Basic Information

* Process Name: repe

Description: east

Intent: Unidirectional

* Credential Domain: OracleGoldenGate

* Credential Alias: ggwest

Source: Trail

* Trail Name: ea

Trail Subdirectory: east

Begin: Position in Log

* Transaction Log Sequence Number 0

* Transaction Log RBA Offset 0

Checkpoint Table: "DBWEST"."GGADMIN"."GGS_C...

4. Click **Next** to view the **Replicat Parameter File** screen. All the parameters that you have specified are available for review here.
   For multitenant container databases, Replicat can only apply to one pluggable database. To specify the correct one, use a SQL*Net connect string for the database user that you specify with the `USERID` or `USERIDALIAS` parameter. For example: `ggadmin@DBWEST`. In the parameter file, specify only the *`schema.object`* in the `TARGET` portion of the `MAP` statements. In the `MAP` portion, identify source objects captured from more than one pluggable database with their three-part names or use the `SOURCECATALOG` parameter with two-part names.

   Here's a sample of the Replicat Parameter File:

```
REPLICAT repe
USERIDALIAS ggwest DOMAIN OracleGoldenGate
--DDL EXCLUDE ALL
DDLERROR default discard
REPERROR (default,discard)
DDLOPTIONS REPORT
SOURCECATALOG DBEAST
MAP hr.*, TARGET hr.*;
```

After the Replicat starts successfully, you can see the Extract and Replicat processes in running state on the **Administration Service** Overview page.

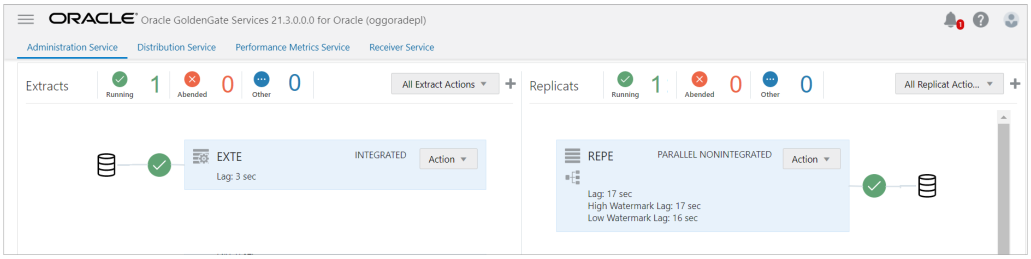

Extracts Running 1 Abended 0 Other 0

Replicats Running 1 Abended 0 Other 0

EXTE INTEGRATED
Lag: 3 sec

REPE PARALLEL NONINTEGRATED
Lag: 17 sec
High Watermark Lag: 17 sec
Low Watermark Lag: 16 sec

Step 5: Test the Replication

To test if the replication has started, try insert, update, or delete operations on your database and then follow these steps:

1. Click **Action** from the Extract (**exte**) section and click **Details**.
2. Click the **Statistics** tab. You'll see additions to the **Insert**, **Updates**, or **Deletes** columns on this page.

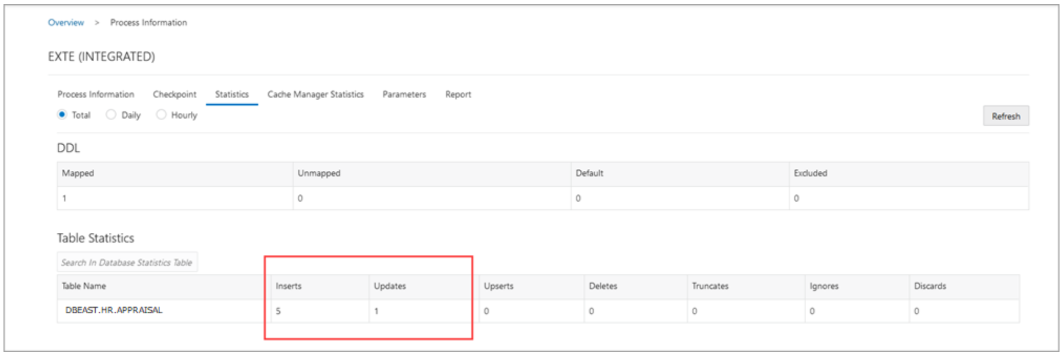

EXTE (INTEGRATED)

● Total ○ Daily ○ Hourly

DDL

Mapped	Unmapped	Default	Excluded
1	0	0	0

Table Statistics

Table Name	Inserts	Updates	Upserts	Deletes	Truncates	Ignores	Discards
DBEAST.HR.APPRAISAL	5	1	0	0	0	0	0

Also see the **Statistics** tab using the Replicat **Details** option. You would see the updates in the **Table Statistics** section.

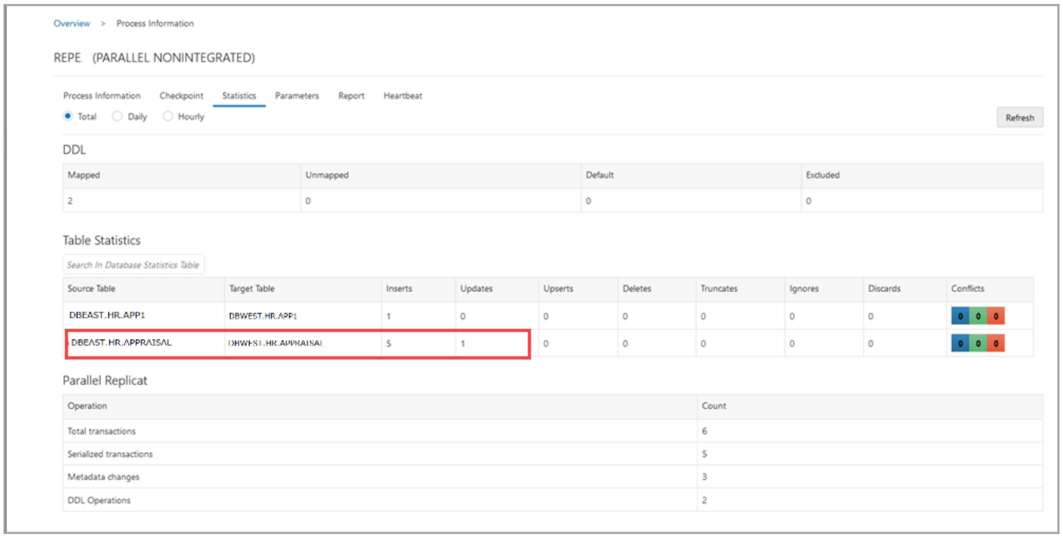

REPE (PARALLEL NONINTEGRATED)

● Total ○ Daily ○ Hourly

DDL

Mapped	Unmapped	Default	Excluded
2	0	0	0

Table Statistics

Source Table	Target Table	Inserts	Updates	Upserts	Deletes	Truncates	Ignores	Discards	Conflicts
DBEAST.HR.APP1	DBWEST.HR.APP1	1	0	0	0	0	0	0	0 0 0
DBEAST.HR.APPRAISAL	DBWEST.HR.APPRAISAL	5	1	0	0	0	0	0	0 0 0

Parallel Replicat

Operation	Count
Total transactions	6
Serialized transactions	5
Metadata changes	3
DDL Operations	2

1.2 Set Up Bidirectional Replication for Oracle GoldenGate Microservices Architecture

This quickstart demonstrates an active-active bidirectional replication between two pluggable databases over a single multitenant container Oracle database instance.

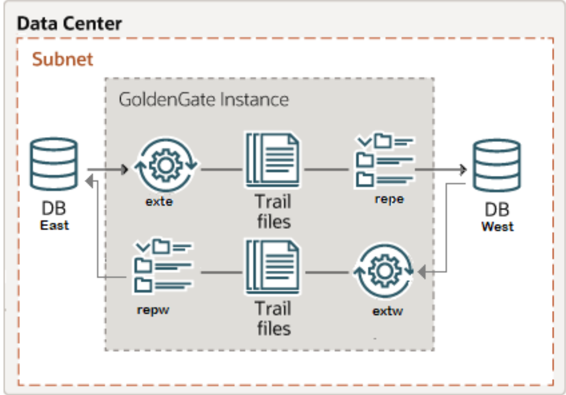

An active-active bidirectional replication implies that both data sources and targets (PDBs in this case), have the potential to send updates to each other. There are two data sources with identical sets of data that can be changed by application users on either side. Oracle GoldenGate replicates transactional data changes from each database to the other to keep both sets of data current.

The following diagram depicts the bidirectional replication workflow shown in this quickstart:

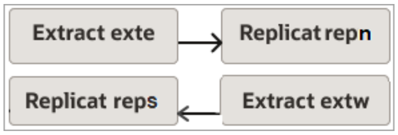

Note:

This quickstart uses a single multitenant container database with two PDBs to demonstrate bidirectional replication between two PDBs. However, in most real-life scenarios, bidirectional data replication happens across different multitenant container databases or different database instances.

Here's what's covered:

- Process Names in the Bidirectional Data Replication Environment
- Considerations for Configuring a Bidirectional Replication
- Oracle Multitenant Container Database Privileges Required for Data Replication
- Automatic Conflict Detection and Resolution (ACDR) Configuration
- Bidirectional Data Replication Process Configuration

Process Names in the Bidirectional Data Replication Environment

The following nomenclature is used to refer to processes for the database and Oracle GoldenGate.

Container Database (CDB$ROOT) Process Names	Pluggable Database (DBEAST) Process Names	Pluggable Database (DBWEST) Process Name
• **CDB$ROOT database user: c##ggadmin** • **Database credential alias: cggnorth**	• Database user: **ggadmin** • Database alias: **ggeast** • Extract: **exte** • Replicat **repn**	• Database user: **ggadmin** • Database alias: **ggwest** • Extract: **extw** • Replicat: **reps**

On DBWest:

Considerations for Configuring a Bidirectional Replication

To maintain data integrity and avoid conflicts, you need to configure the Extract and Replicat processes to prevent data looping and conflict using certain parameters and the automatic conflict detection and resolution (ACDR) feature.

Ideally, all situations that could lead to potential conflicts in a bidirectional or multidirectional replication must be avoided. However, if conflicts occur, Oracle GoldenGate provides the automatic conflict detection and resoution (ACDR) feature to handle them.

- **At the PDB level:**

  The Automatic Conflict Detection and Resolution feature (ACDR) available with Oracle database, allows you to manage conflict detection and resolution using the DBMS_GOLDENGATE_ADM package, using the ADD_AUTO_CDR procedure. You need to enable this package at the database level on both PDBs in this case. See Enable ACDR.

- **Oracle GoldenGate Extract parameter settings**

  – LOGALLSUPCOLS: This parameter controls writing of supplementally logged columns specified using ADD TRANDATA and the columns enabled for Conflict Detection and Resolution (CDR) in Oracle GoldenGate. This parameter is set by default for Extract.

  – UPDATERECORDFORMAT: This parameter is set by default for integrated Extract, so don't need to set it in the parameter file. Its function is to combine the before and after images of an UPDATE operation into a single record in the trail. The COMPACT option generates one trail record that contains the before and after images of an UPDATE, where the before image includes all the columns that are available in the transaction record, but the after image is limited to the primary key columns and the columns that were modified in the UPDATE.

  – EXCLUDETAG option ensures that there is not looping of data. Looping of data happens when a database sends updates to the second database and the second database assumes those updates to be a new changes, and tries to replicate this update back to the source database itself. These parameter settings are done when configuring the Extract parameter file, as shown in Step 3: Add Extracts of this document.

- **Oracle GoldenGate Replicat parameter settings**:

  ACDR works with integrated Replicat or parallel integrated Replicat. See the Replicat Parameter file in this document to know more.

Set the Required Privileges for Oracle Multitenant Database

In Oracle database, you need to enable replication for Oracle GoldenGate and assign privileges to the database user at the CDB level and the pluggable database (PDB) level.

The database is in ARCHIVELOG mode and FORCE LOGGING and Supplemental Logging is enabled. For the container database, assign the following privileges to the common user (cdb$root):

CDB User Privileges

```
## CGGNORTH DATABASE SETUP AT CDB LEVEL
alter session set container=cdb$root;
alter system set enable_goldengate_replication=TRUE;
alter system set streams_pool_size=2G;
alter database force logging;
alter database add supplemental log data;
create user c##ggadmin identified by Password container=all default
tablespace GG_DATA temporary tablespace temp;
grant connect, resource, dba to c##ggadmin container=all;
grant create session to c##ggadmin container=all;
exec
dbms_goldengate_auth.grant_admin_privilege('c##ggadmin',container=>'all
');
```

PDB User Privileges for DBEAST

```
alter session set container=DBEAST;
create user ggadmin identified by Password container=current;
grant connect, resource, dba to ggadmin container=current;
grant create session to ggadmin container=current;
exec dbms_goldengate_auth.grant_admin_privilege('ggadmin');
```

PDB User Privileges for DBWEST

```
alter session set container=DBWEST;
create user ggadmin identified by password container=current;
grant connect, resource, dba to ggadmin container=current;
grant create session to ggadmin container=current;
exec dbms_goldengate_auth.grant_admin_privilege('ggadmin');
```

Note:

Granting DBA role is not mandatory for every user. Privileges should be granted depending on the actions that the user needs to perform on the database. For example, to grant DML operation privileges to insert, update, and delete transactions to ggadmin, use the GRANT ANY INSERT/UPDATE/DELETE privileges and to further allow users to work with tables and indexes as part of DML operations, use the GRANT CREATE/DROP/ALTER ANY TABLE/INDEX privileges. In this quickstart, the assumption is that the database user is a database administrator. See Oracle Database Privileges and Other Requirements for Multitenant Container Databases to know more about specific privilege requirements.

Enable ACDR

Before enabling ACDR at the database level, it is recommended that you stop any running Extract or Replicat processes. To enable ACDR for the PDB DBEAST:

```
exec dbms_goldengate_adm.add_auto_cdr('hr', 'employees',
record_conflicts=>TRUE);
```

The output will show as:

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

Now, switch to the other PDB, DBWEST and run the same command:

```
exec dbms_goldengate_adm.add_auto_cdr('hr', 'employees');
```

This enables the ACDR package on both PDBs.

You can check if ACDR has been enabled for the PDBs by checking for invisible columns that are added to manage ACDR at the column level. Run the following commands to test this:

Use the view all_gg_auto_cdr_tables to list down the columns used for ACDR in the PDBs:

```
SELECT table_owner, table_name, tombstone_table, row_resolution_column, FROM
all_gg_auto_cdr_tables;
```

The output for this command shows:

```
TABLE_OWNER
-----------------------------------------------------------------------------
---
TABLE_NAME
-----------------------------------------------------------------------------
---
TOMBSTONE_TABLE
-----------------------------------------------------------------------------
---
ROW_RESOLUTION_COLUMN
-----------------------------------------------------------------------------
---
HR
EMPLOYEES
DT$_EMPLOYEES
CDRTS$ROW
```

Notice the two invisible columns that are added here:

- DT$_EMPLOYEES: This is the tombstone table used for locking any delete transactions.
- CDRTS$ROW: This is the row resolution column. When there is a conflict, this column which contains the timestamp for the transaction, is used to decide the record that would be

applied in a row. This implies that the record with the latest timestamp would be used to apply changes in the row.

These columns are appended to `schema.table` on both PDBs, `DBEAST` and `DBWEST`.

After you have enabled ACDR, you'll need to edit the Replicat parameter file to include the invisible columns. Add the `MAPINVISIBLECOLUMNS` parameter in the Replicat parameter file, to allow Replicat to include target columns with default column mapping. This is explained in detail when configuring the Replicat Parameter File in **Step 4: Add Replicat** section.

Restart the Extract and Replicat processes from the web interface:

1. Log in to the Administration Service web interface.
2. From the Administration Service Overview page, click the **Action** button next to the Extract process, **exte**.
3. Click **Start**.

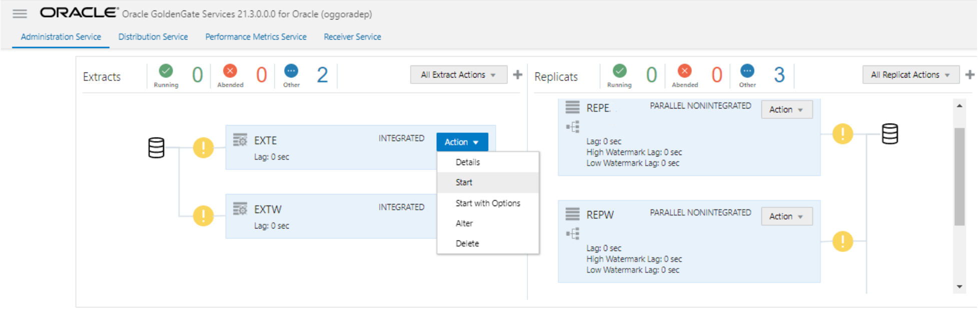

The green check mark would appear next to the process indicating that the processes started successfully.

Similarly, start the other Extract and Replicat processes on both PDBs. To know more about different strategies for ACDR, see Configuring Automatic Conflict Detection and Resolution.

Configure the Replication Process from Oracle GoldenGate MA Web Interface

Using the following steps, you'll be able to configure data capture (Extract) and apply (Replicat) processes. You'll also be able to test if the replication has started.

- Step 1: Add Database Credentials from the Administration Service
- Step 2: Add Heartbeat and Checkpoint Tables
- Step 3: Add Extracts
- Step 4: Add a Replicat
- Test and Monitor Transactions
- Test Automatic Conflict Detection and Resolution

The `DISTPATH` process is not used for this configuration.

Step 1: Add Database Credentials from the Administration Service

In this section, you'll add the database credentials to connect to the source and target databases using EZConnect.

1. Keep your database user credentials, which created in the previous session, ready. You'll use them to connect Oracle GoldenGate to the database server.
2. Open the Service Manager login page in a web browser and log in to the Service Manager with your Oracle GoldenGate administrator user credentials. If logging in for the first time, you have to log in with the administrator account user credentials, created when adding your deployment with Oracle GoldenGate Configuration Assistant wizard.
3. From the Service Manager Overview page, click the port number for the Administration Service of the deployment.

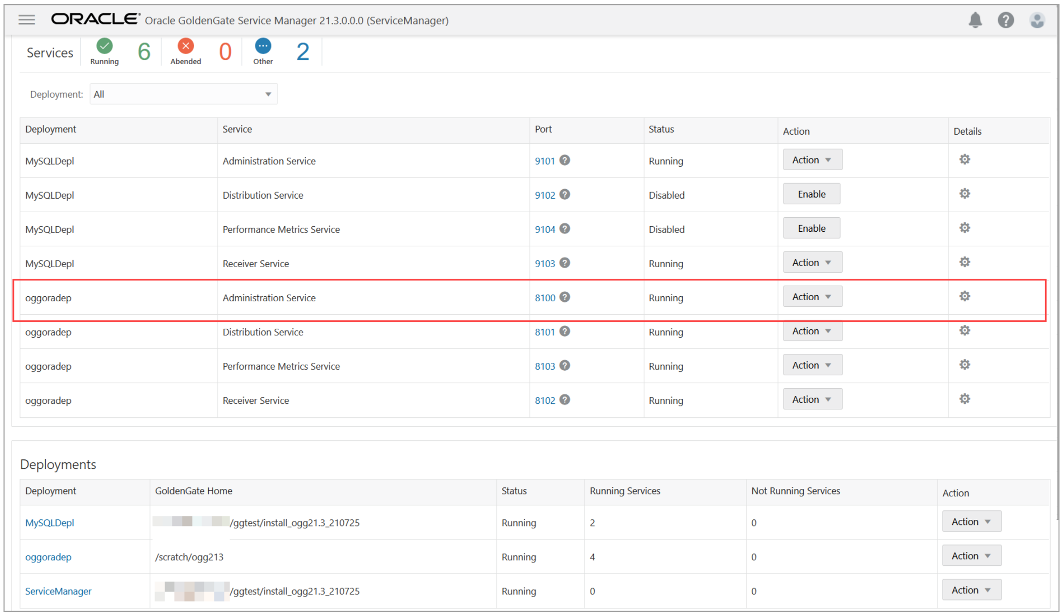
ORACLE Oracle GoldenGate Service Manager 21.3.0.0.0 (ServiceManager)
Services
Running 6
Abended 0
Other 2
Deployment: All

Deployment	Service	Port	Status	Action	Details
MySQLDepl	Administration Service	9101	Running	Action	⚙
MySQLDepl	Distribution Service	9102	Disabled	Enable	⚙
MySQLDepl	Performance Metrics Service	9104	Disabled	Enable	⚙
MySQLDepl	Receiver Service	9103	Running	Action	⚙
oggoradep	Administration Service	8100	Running	Action	⚙
oggoradep	Distribution Service	8101	Running	Action	⚙
oggoradep	Performance Metrics Service	8103	Running	Action	⚙
oggoradep	Receiver Service	8102	Running	Action	⚙

Deployments

Deployment	GoldenGate Home	Status	Running Services	Not Running Services	Action
MySQLDepl	[redacted]/ggtest/install_ogg21.3_210725	Running	2	0	Action
oggoradep	/scratch/ogg213	Running	4	0	Action
ServiceManager	[redacted]/ggtest/install_ogg21.3_210725	Running	0	0	Action

This opens the Administration Service login page.

4. Log in to the Administration Service using the same credentials, which you used to log in to the Service Manager. The Administration Service **Overview** page is displayed.

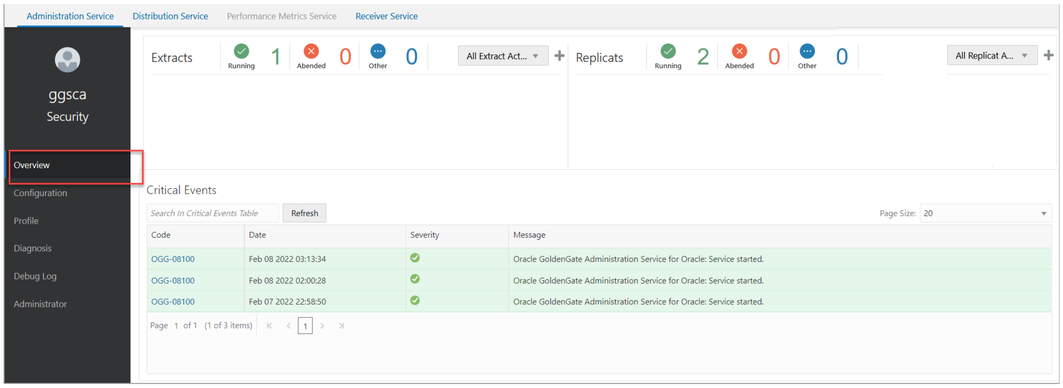

Code	Date	Severity	Message
OGG-08100	Feb 08 2022 03:13:34	✓	Oracle GoldenGate Administration Service for Oracle: Service started.
OGG-08100	Feb 08 2022 02:00:28	✓	Oracle GoldenGate Administration Service for Oracle: Service started.
OGG-08100	Feb 07 2022 22:58:50	✓	Oracle GoldenGate Administration Service for Oracle: Service started.

5. Click the **Application Navigation** icon to open the left-navigation pane and click **Configuration** to open the **Database** tab of the **Configuration** page.

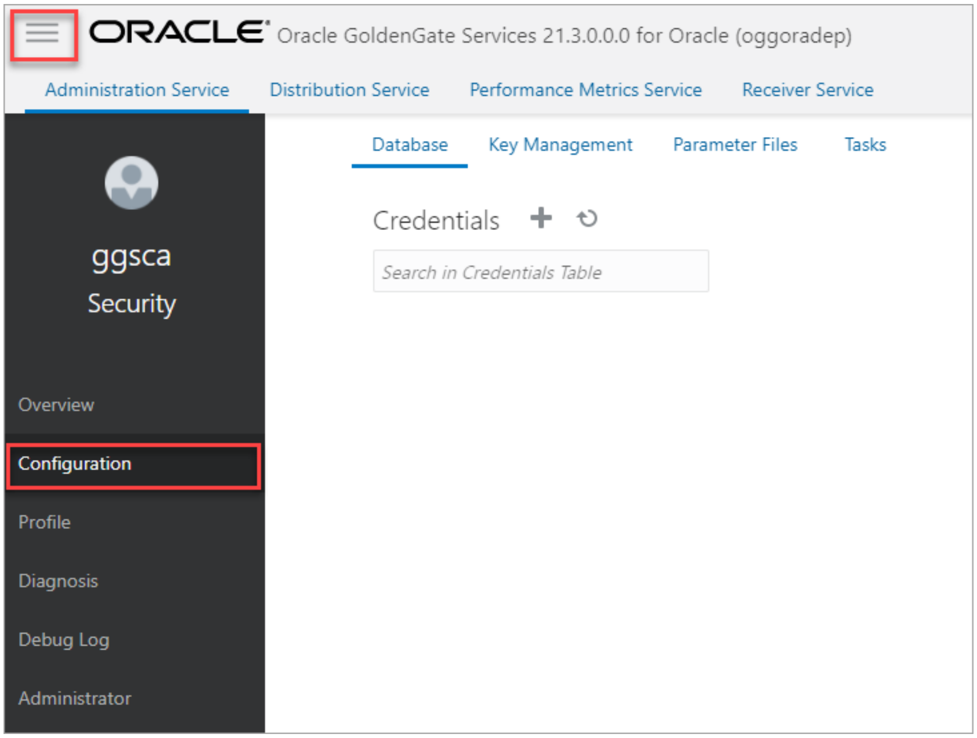

6. Click the plus (+) sign in the **Credentials** section to begin adding database user credentials.

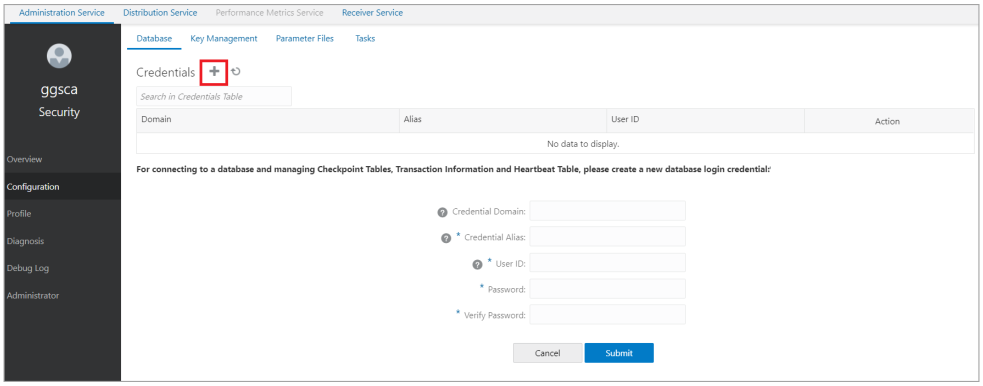

Domain	Alias	User ID	Action
No data to display.			

For connecting to a database and managing Checkpoint Tables, Transaction Information and Heartbeat Table, please create a new database login credential:

Credential Domain:
* Credential Alias:
* User ID:
* Password:
* Verify Password:

7. You need to add connections for container database (CDB) and pluggable databases (PDBs). Each CDB is used to capture (Extract) from the source database and PDB for delivery (Replicat).

   Use the EZconnect syntax to configure the database connection. You need the username, password, hostname, port number, and service name connection information to use the EZConnect syntax.

   Here's the syntax that you need to specify in the User ID field:

   ```
   username@hostname:port/service_name
   ```

   Here's an example for setting the User ID with EZConnect:

   ```
   c##ggadmin@dc.example.com:1521/DBWEST.example.com
   ```

8. Click on the blue icon in the **Actions** column to connect to the database. The icon turns blue when the connection is successful.

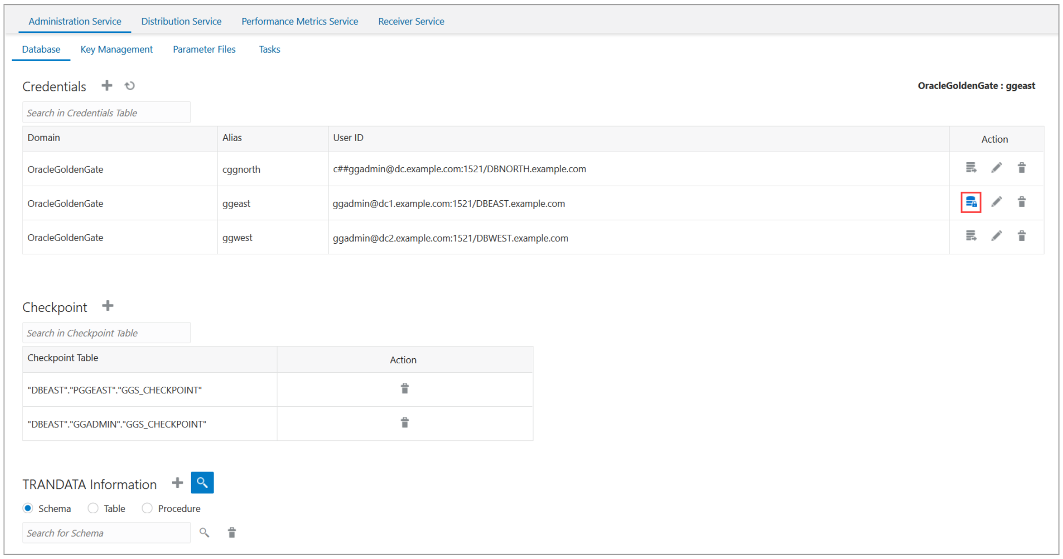

Domain	Alias	User ID	Action
OracleGoldenGate	cggnorth	c##ggadmin@dc.example.com:1521/DBNORTH.example.com	
OracleGoldenGate	ggeast	ggadmin@dc1.example.com:1521/DBEAST.example.com	
OracleGoldenGate	ggwest	ggadmin@dc2.example.com:1521/DBWEST.example.com	

Checkpoint Table	Action
"DBEAST"."PGGEAST"."GGS_CHECKPOINT"	
"DBEAST"."GGADMIN"."GGS_CHECKPOINT"	

After connecting to the database, the sections to add checkpoint table, TRANDATA, and heartbeat table are displayed.

Step 2: Add Heartbeat, and Checkpoint Tables

Add the heartbeat tables for the PDBs to monitor any possible lags. Add a checkpoint table for the target database to ensure that if there is a failure, then the Extract and Replicat processes can restart from the point of failure.

> **Note:**
>
> You don't need to add TRANDATA as this is internally done with the PL/SQL call of `ADD_AUTO_CDR`. You might want to check that supplemental logging is enabled for the tables.

1. Use the **TRANDATA Information** section to check if supplemental logging has been enabled for the tables set up for capture.

   You can search for the schema for which you added the trandata, using the magnifier glass search icon. This will display the trandata information. The following image shows the trandata information for the `HR` schema in the pluggable database `DBEAST`.

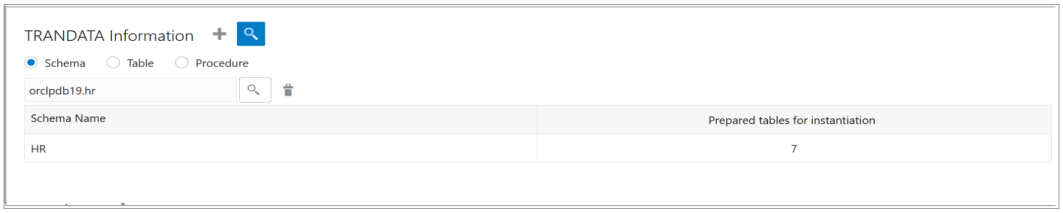

orclpdb19.hr

Schema Name	Prepared tables for instantiation
HR	7

See Configuring Logging Properties to learn the steps for configuring the logging properties at the Schema, Table or Procedure level.

2. To set up the checkpoint table for Replicat, you need to connect to the target database credentials (**ggwest**) from the Credentials section.

3. Click the plus sign (+) to add the checkpoint table for the PDBs.

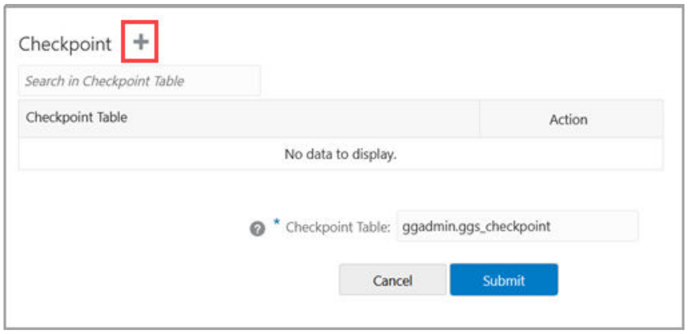

Checkpoint Table	Action
No data to display.	

* Checkpoint Table: ggadmin.ggs_checkpoint

Click **Submit**. The checkpoint table is added.

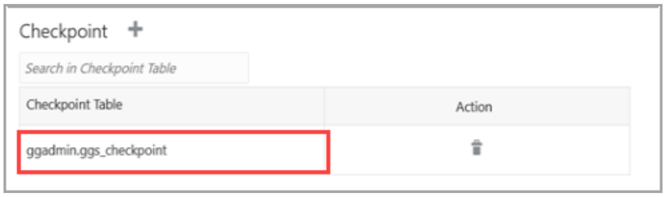

Checkpoint Table	Action
ggadmin.ggs_checkpoint	🗑

Also see the Before Creating an Extract section, for details on creating heartbeat tables.

4. Add another checkpoint table for the second Replicat, `reps`, by repeating the steps 3 and 4.

5. Add the heartbeat tables for both source and target endpoints by connecting to **ggeast** and **ggwest** database credential aliases. Add the heartbeat table by clicking the plus sign.

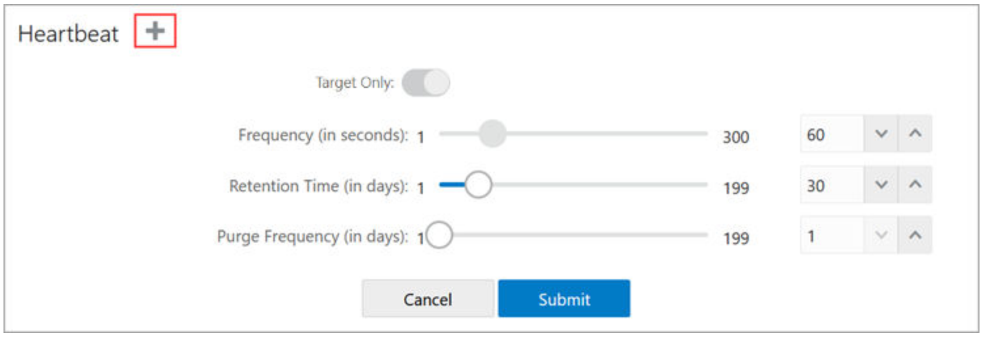

Target Only: ○

Frequency (in seconds): 1 — 300: 60

Retention Time (in days): 1 — 199: 30

Purge Frequency (in days): 1 — 199: 1

6. Click **Submit** after adjusting the heartbeat options.

Step 3: Add Extracts

In this section, you will add two extracts, **`exte`** and **`extw`**. The Extract process captures data from the source database and writes it to a trail file. The trail file for **`exte`** is **`ea`** and for **`extw`** is **`ew`**.

1. Click the **Overview** option from the left-navigation pane of the Administration Service and click the plus sign (+) from the Extract section.

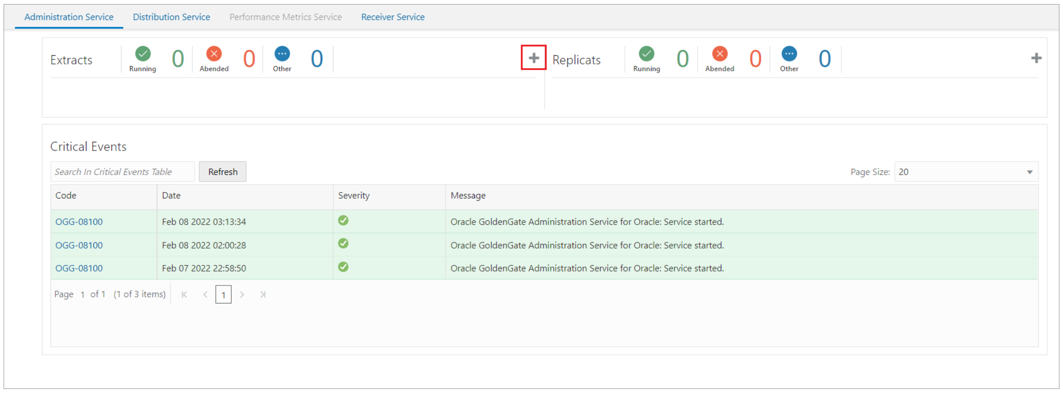

Code	Date	Severity	Message
OGG-08100	Feb 08 2022 03:13:34	✓	Oracle GoldenGate Administration Service for Oracle: Service started.
OGG-08100	Feb 08 2022 02:00:28	✓	Oracle GoldenGate Administration Service for Oracle: Service started.
OGG-08100	Feb 07 2022 22:58:50	✓	Oracle GoldenGate Administration Service for Oracle: Service started.

2. From Add Extract wizard, select **Integrated Extract**.

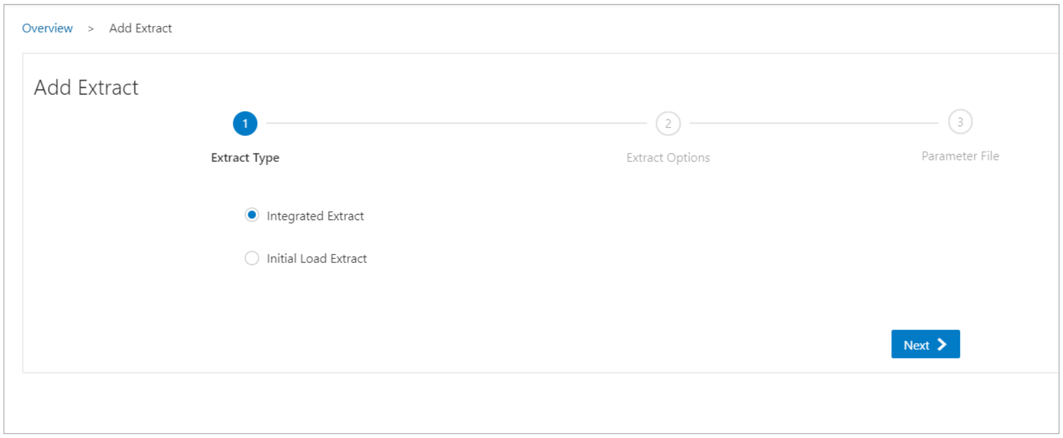

● Integrated Extract

○ Initial Load Extract

3. Click **Next** and specify the Extract options in the Extract Options screen. See the detailed steps to add an Extract from the Add an Extract section.

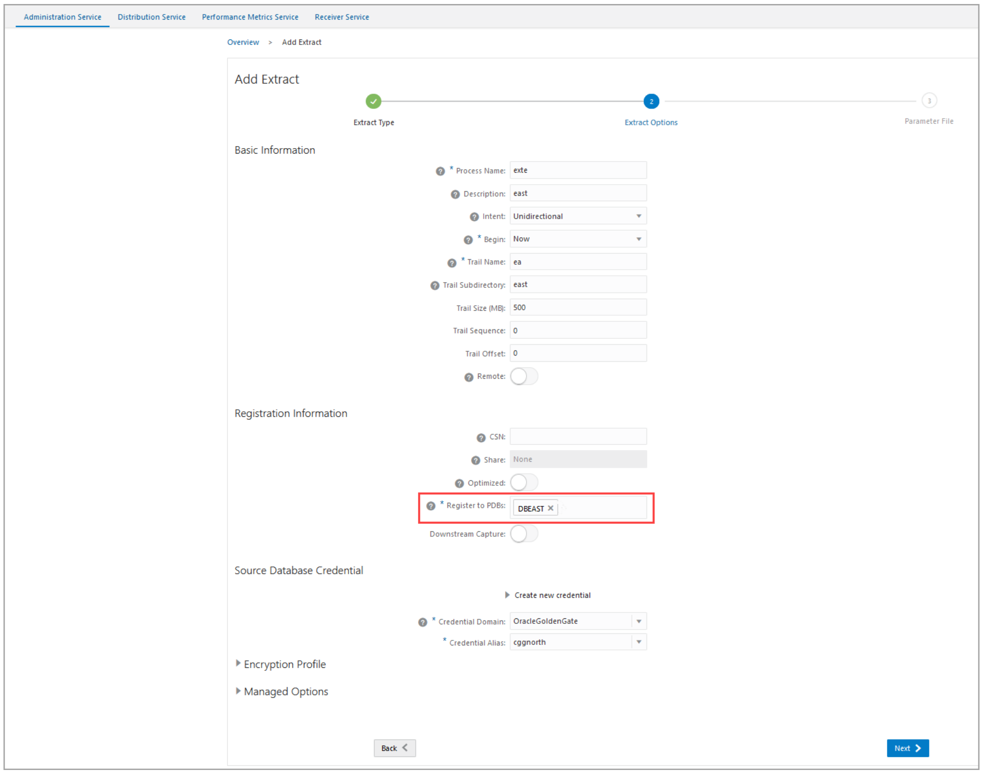

Add Extract

Basic Information

* Process Name: exte
Description: east
Intent: Unidirectional
* Begin: Now
* Trail Name: ea
Trail Subdirectory: east
Trail Size (MB): 500
Trail Sequence: 0
Trail Offset: 0
Remote: ○

Registration Information

CSN:
Share: None
Optimized: ○
* Register to PDBs: DBEAST ✕
Downstream Capture: ○

Source Database Credential

▸ Create new credential
* Credential Domain: OracleGoldenGate
* Credential Alias: cggnorth

▸ Encryption Profile

▸ Managed Options

If you are creating the Extract for a pluggable database, then you'll see option **Register to PDBs** as soon as you enter the credentials domain and alias. Select the PDB in the container database that you want to use for replication.

4. After you enter the options for the Extract (`exte`), click **Next**. The next screen displays the Extract parameter file to help you review the Extract settings.

   Here's the Extract parameter file for the Extract exte:

   ```
   EXTRACT exte
   USERIDALIAS cggnorth DOMAIN OracleGoldenGate
   EXTTRAIL east/ea
   SOURCECATALOG DBEAST
   TRANLOGOPTIONS EXCLUDETAG 00
   DDL INCLUDE MAPPED OBJNAME hr.*
   DDLOPTIONS REPORT
   TABLE DBEAST.hr.*;
   ```

   Review these settings and update the Extract configuration as needed.

   For multitenant databases, you need to add entries for Extract to capture from multiple pluggable databases to a single trail. In the parameter file, source objects must be specified in `TABLE` statements with the fully qualified three-part names in the format of *`container.schema.object`* or using the `SOURCECATALOG` parameter with two-part names *`schema.object`*.

   Click **Create and Run** to start your Extract.

To create the Extract **`extw`**:

1. Navigate back to the Overview page using the Application Navigation pane.

2. From Add Extract wizard, select Integrated Extract.
3. Click **Next** and specify the Extract options in the Extract Options screen.

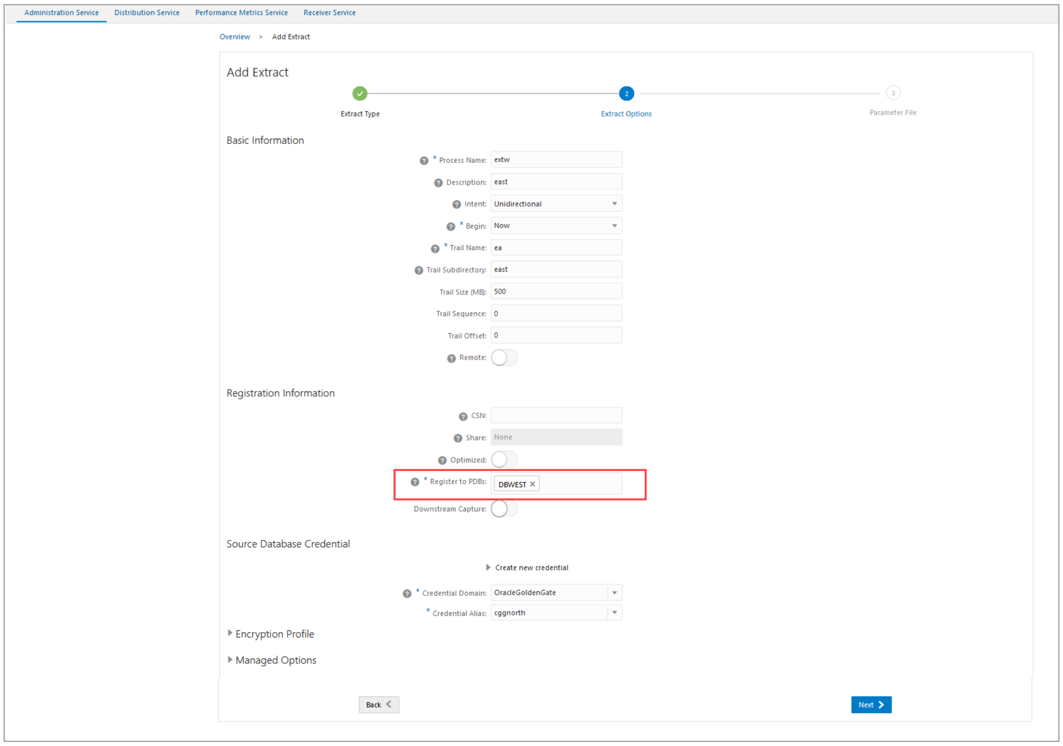

Add Extract

Basic Information

* Process Name: extw
Description: east
Intent: Unidirectional
* Begin: Now
* Trail Name: ea
Trail Subdirectory: east
Trail Size (MB): 500
Trail Sequence: 0
Trail Offset: 0
Remote:

Registration Information

CSN:
Share: None
Optimized:
* Register to PDBs: DBWEST ×
Downstream Capture:

Source Database Credential

Create new credential
* Credential Domain: OracleGoldenGate
* Credential Alias: cggnorth

Encryption Profile

Managed Options

4. Select the PDB as `DBWEST` in the container database that you want to use for replication.
5. After you enter the options for the Extract, click Next. The next screen displays the Extract parameter file to help you review the Extract settings.
6. Enter the options for Extract parameter:

   ```
   EXTRACT extw
   USERIDALIAS cggnorth DOMAIN OracleGoldenGate
   EXTTRAIL west/ew
   SOURCECATALOG DBWEST
   TRANLOGOPTIONS EXCLUDETAG 00
   DDL INCLUDE MAPPED OBJNAME hr.*
   DDLOPTIONS REPORT
   TABLE DBWEST.hr.*;
   ```

   Review these settings and update the Extract configuration as needed.
7. Click **Create and Run** to start your Extract.

Step 4: Add a Replicat

In this section, you will add Replicats **repe** and **repw**. The Replicat process delivers the change data from the trail file (**`ea`**) created by the Extract, to the target database. Replicat reads the trail file on the target database, reconstructs the DML or DDL operations, and applies them to the target database.

1. Before Creating a Replicat, make sure that you added your checkpoint table for the target database (`DBWEST`) by connecting to the **ggwest** database credentials.

2. Select a Replicat type to deliver data to the target database. Follow the wizard to complete adding a Replicat.

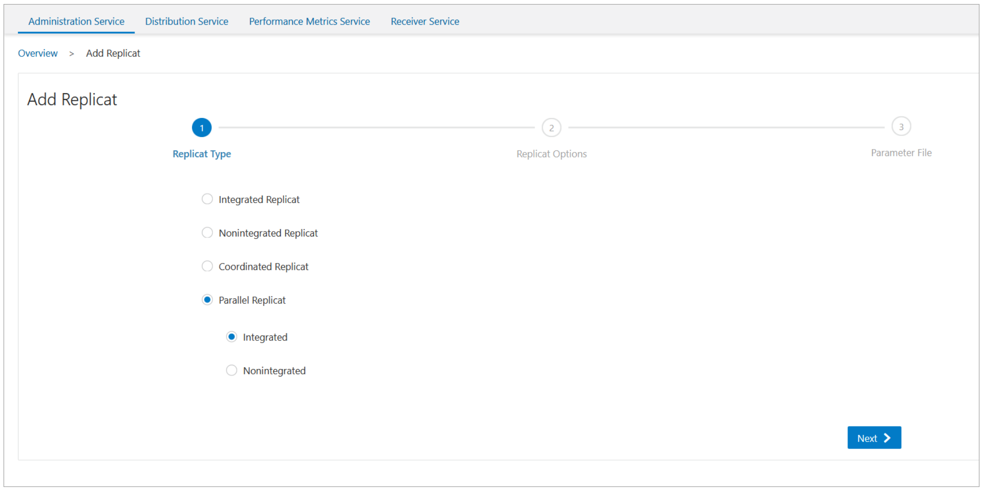

Add Replicat

○ Integrated Replicat

○ Nonintegrated Replicat

○ Coordinated Replicat

● Parallel Replicat

- ● Integrated
- ○ Nonintegrated

3. Select the **Parallel Integrated Replicat** option in the Replicat Options screen.

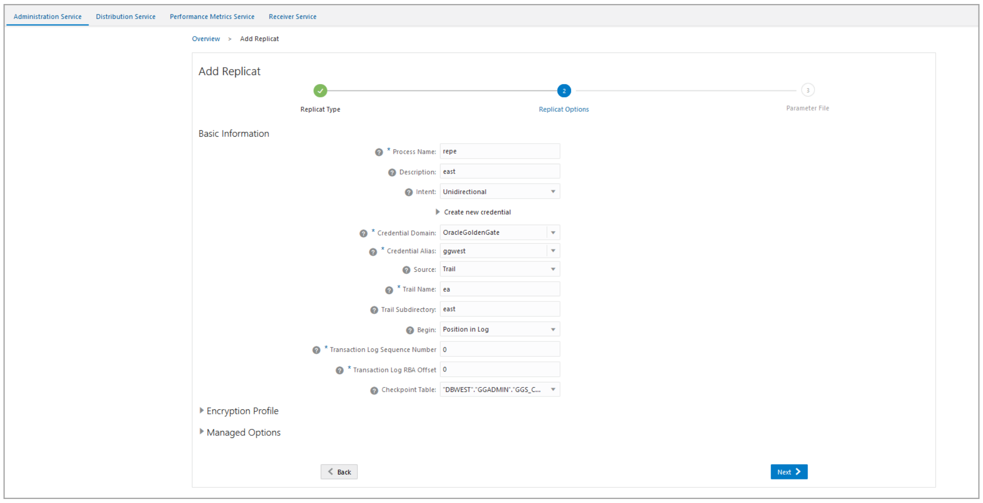

Add Replicat

Basic Information

* Process Name:	repe
Description:	east
Intent:	Unidirectional
Create new credential	
* Credential Domain:	OracleGoldenGate
* Credential Alias:	ggwest
Source:	Trail
* Trail Name:	ea
Trail Subdirectory:	east
Begin:	Position in Log
* Transaction Log Sequence Number	0
* Transaction Log RBA Offset	0
Checkpoint Table:	"DBWEST"."GGADMIN"."GGS_C...

Encryption Profile

Managed Options

4. Click **Next** to view the **Replicat Parameter File** screen. All the parameters that you have specified are available for review here.
For multitenant container databases, Replicat can only apply to one pluggable database. To specify the correct one, use a SQL*Net connect string for the database user that you specify with the `USERID` or `USERIDALIAS` parameter. For example: `ggadmin@DBWEST`.

   In the parameter file, specify only the *`schema.object`* in the `TARGET` portion of the `MAP` statements. In the `MAP` portion, identify source objects captured from more than one pluggable database with their three-part names or use the `SOURCECATALOG` parameter with two-part names.

   In case of integrated parallel Replicat, `MAPINVISIBLECOLUMNS` parameter is set by default. You don't need to set it in the Replicat parameter file explicitely.

Here's a sample of the Replicat Parameter File:

```
REPLICAT repe
USERIDALIAS ggwest DOMAIN OracleGoldenGate
DDLOPTIONS REPORT
SOURCECATALOG DBEAST
MAP hr.*, TARGET hr.*;
```

To create the second Replicat **repw**, follow these steps:

1. Repeat steps 1 and 2 from the steps to add the first Replicat (**`repe`**).
2. In the Replicat options screen, enter the following details:

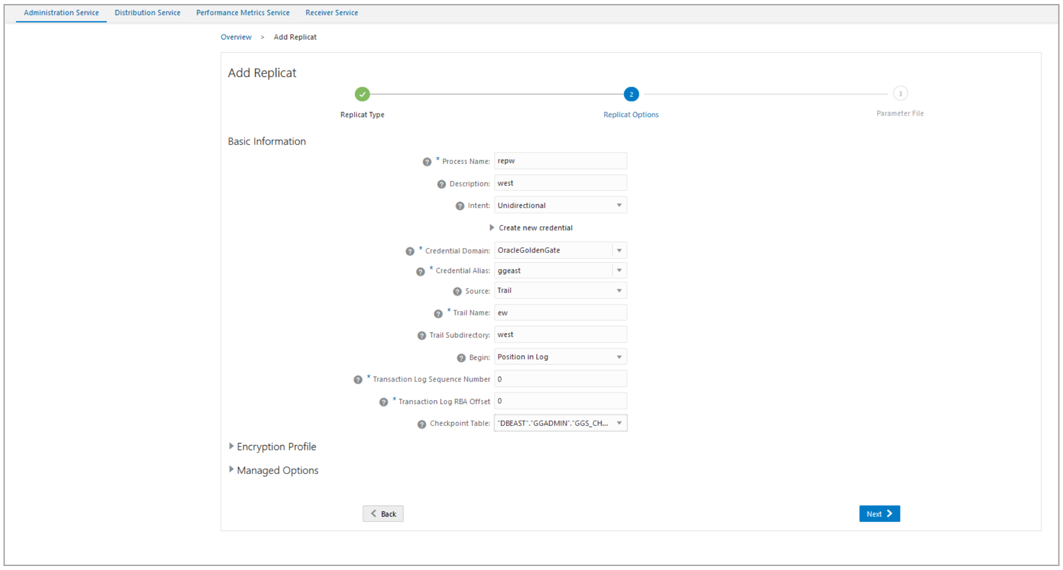

Basic Information

- * Process Name: repw
- Description: west
- Intent: Unidirectional
- * Credential Domain: OracleGoldenGate
- * Credential Alias: ggeast
- Source: Trail
- * Trail Name: ew
- Trail Subdirectory: west
- Begin: Position in Log
- * Transaction Log Sequence Number 0
- * Transaction Log RBA Offset 0
- Checkpoint Table: "DBEAST"."GGADMIN"."GGS_CH...

Apart from entering the other options, make sure you enter the following details:

a. Specify the trail name as ew and the trail file subdirectory as west.

b. Select the checkpoint table as `DBWEST.ggs_checkpoint`.

c. Click Next.

d. Change or modify the Replicat parameter file, as follows:

```
REPLICAT repw
USERIDALIAS ggeast DOMAIN OracleGoldenGate
SOURCECATALOG DBWEST
DDL INCLUDE ALL
DDLOPTIONS REPORT
MAPEXCLUDE ggadmin.ggs_checkpoint*
MAPINVISIBLECOLUMNS
MAP hr.*, TARGET hr.*;
```

After the Replicat starts successfully, you can see the Extract and Replicat processes in running state on the **Administration Service** Overview page.

Test and Monitor Transactions

The following screen shows that records were captured by the **`exte`** Extract from the `hr.employees` table on `DBEAST`.

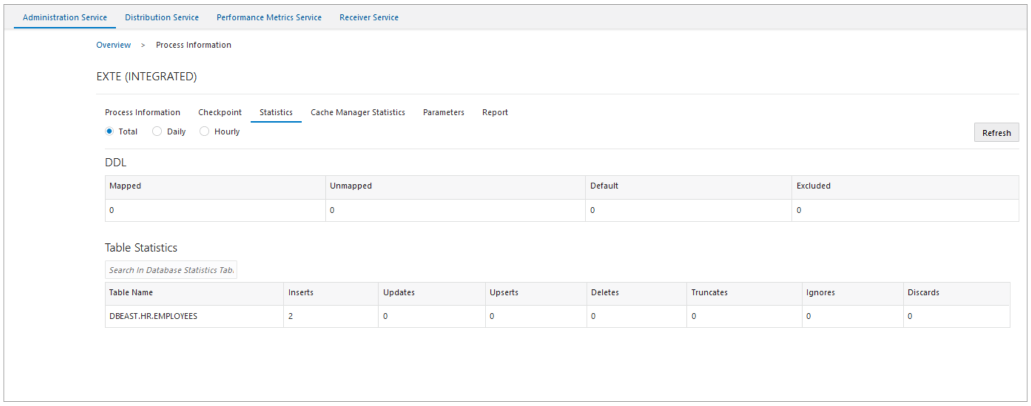

EXTE (INTEGRATED)

● Total ○ Daily ○ Hourly

DDL

Mapped	Unmapped	Default	Excluded
0	0	0	0

Table Statistics

Table Name	Inserts	Updates	Upserts	Deletes	Truncates	Ignores	Discards
DBEAST.HR.EMPLOYEES	2	0	0	0	0	0	0

Check that the same is updated on the Replicat (**`repe`**) as well:

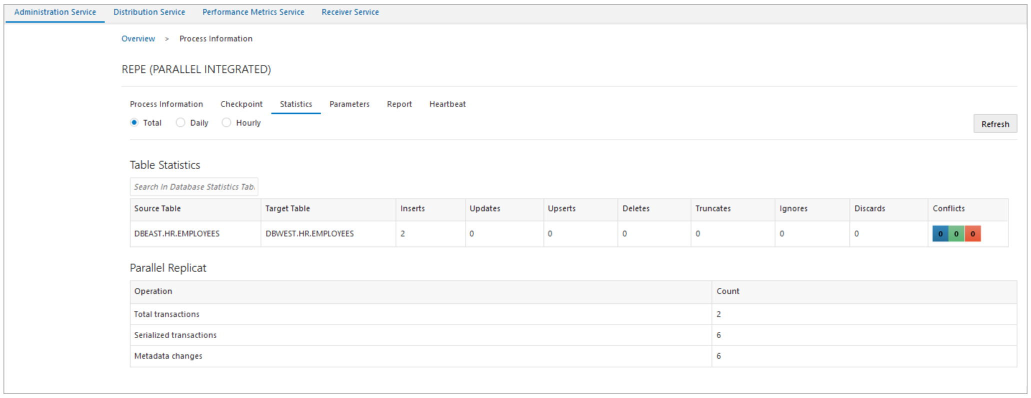

REPE (PARALLEL INTEGRATED)

● Total ○ Daily ○ Hourly

Table Statistics

Source Table	Target Table	Inserts	Updates	Upserts	Deletes	Truncates	Ignores	Discards	Conflicts
DBEAST.HR.EMPLOYEES	DBWEST.HR.EMPLOYEES	2	0	0	0	0	0	0	0 0 0

Parallel Replicat

Operation	Count
Total transactions	2
Serialized transactions	6
Metadata changes	6

The 2 records in the `hr.employees` table are replicated to the endpoint (`DBWEST`).

Let's see the Extract (**`extw`**) on `DBWEST`.

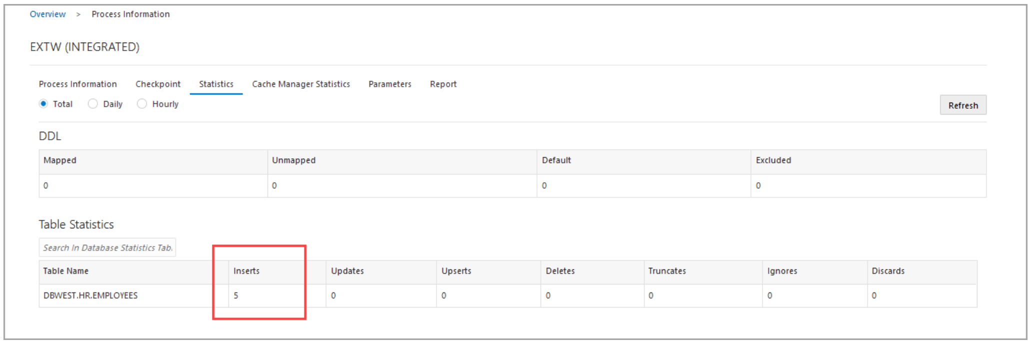

EXTW (INTEGRATED)

● Total ○ Daily ○ Hourly

DDL

Mapped	Unmapped	Default	Excluded
0	0	0	0

Table Statistics

Table Name	Inserts	Updates	Upserts	Deletes	Truncates	Ignores	Discards
DBWEST.HR.EMPLOYEES	5	0	0	0	0	0	0

Notice that the **value of inserted records is 5**. Out of these 5 records, **2 were replicated by `repe`** into `hr.employees` on **`DBWEST`**. **3 new records** were then inserted into `hr.employees` on **`DBWEST`**.

When these 3 records are inserted in the `hr.employees` table in the PDB `DBWEST`, then only the updated records should be replicated in `DBEAST`. The following screen shows that only the updated records are added to `DBEAST`.

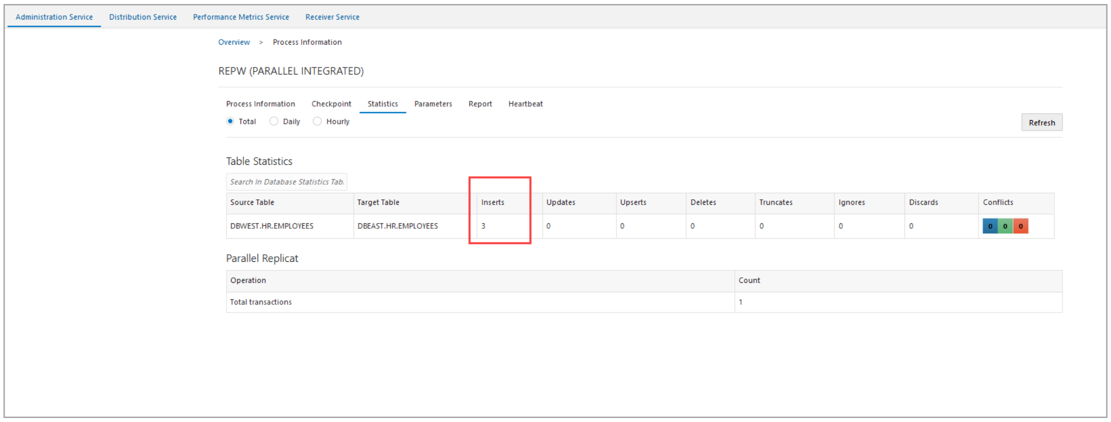

REPW (PARALLEL INTEGRATED)

Table Statistics

Source Table	Target Table	Inserts	Updates	Upserts	Deletes	Truncates	Ignores	Discards	Conflicts
DBWEST.HR.EMPLOYEES	DBEAST.HR.EMPLOYEES	3	0	0	0	0	0	0	0 0 0

Parallel Replicat

Operation	Count
Total transactions	1

As shown in this figure, there are 3 INSERTS, indicating that there was no duplication of records.

This is one way of implementing an active-active bidirectional replication in Oracle GoldenGate MA.

Test Automatic Conflict Detection and Resolution

In this section, the latest timestamp of a record is checked to check if ACDR is able to resolve the conflict in records. To check automatic resolution of conflicts, let's create the following records.

Transaction in `DBEAST`:

In the following example, UPDATE transactions have been run simulateneously on `DBEAST` and `DBWEST` and with ACDR, the conflict is detected and resolved.

Here's the query to update a records in hr.employees on `DBEAST`:

```
UPDATE hr.employees set LAST_NAME='Simmonds', EMAIL='HSIMMONDS' where
EMPLOYEE_ID=204;


UPDATE hr.employees set SALARY='15000' where EMPLOYEE_ID=203;
```

Simulteneously, another query is run on `DBWEST` for the same rows, as shown in the following example:

```
UPDATE hr.employees set LAST_NAME='Symmonds', EMAIL='HSYMMONDS' where
EMPLOYEE_ID=204;


UPDATE hr.employees set SALARY='25000' where EMPLOYEE_ID=203;
```

To check which of these entries was the winner or the entry that was finally applied, and to know the criteria used to apply that entry, use the following options:

Use DBA_APPLY_ERROR_MESSAGES view

On DBEAST, run the following query:

```
select OBJECT_NAME, CONFLICT_TYPE,APPLIED_STATE,CONFLICT_INFO from
DBA_APPLY_ERROR_MESSAGES;
```

The output for this query displays the following:

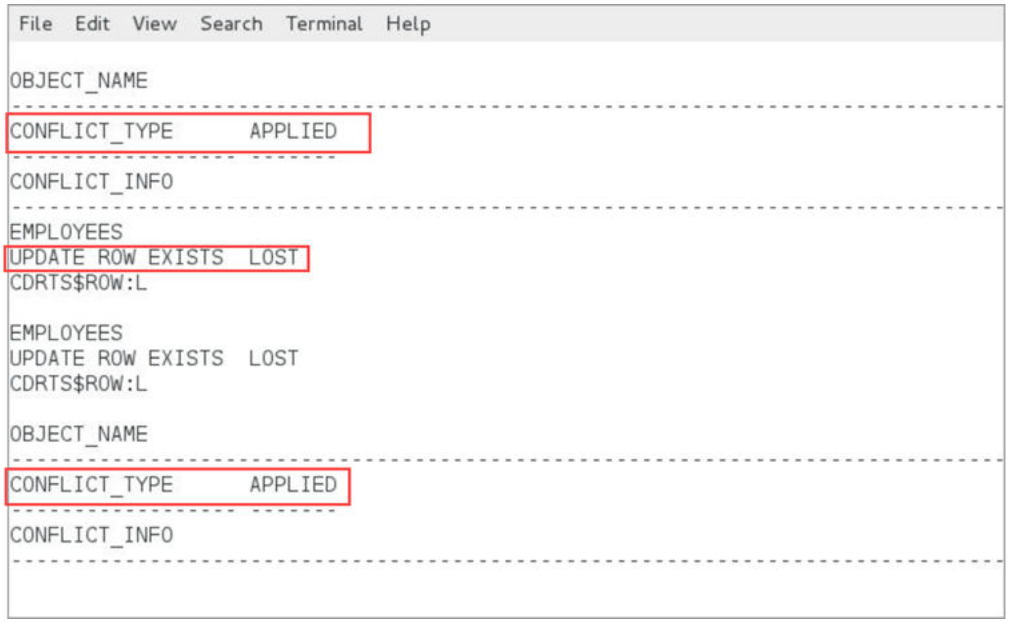

```
OBJECT_NAME
--------------------------------------------------------------------------------
CONFLICT_TYPE       APPLIED
------------------- -------
CONFLICT_INFO
--------------------------------------------------------------------------------
EMPLOYEES
UPDATE ROW EXISTS   LOST
CDRTS$ROW:L

EMPLOYEES
UPDATE ROW EXISTS   LOST
CDRTS$ROW:L

OBJECT_NAME
--------------------------------------------------------------------------------
CONFLICT_TYPE       APPLIED
------------------- -------
CONFLICT_INFO
--------------------------------------------------------------------------------
```

Run the same query on DBWEST also.

Make the CDRTS$ROW visible by running the following command:

```
ALTER TABLE hr.employees modify CDRTS$ROW visible;
```

To check the record that was eventually applied to the rows, run the SELECT query on the table hr.employees. You can run this query on either DBEAST or DBWEST:

```
SELECT * from hr.employees WHERE employee_id=204
```

The output shows as follows:

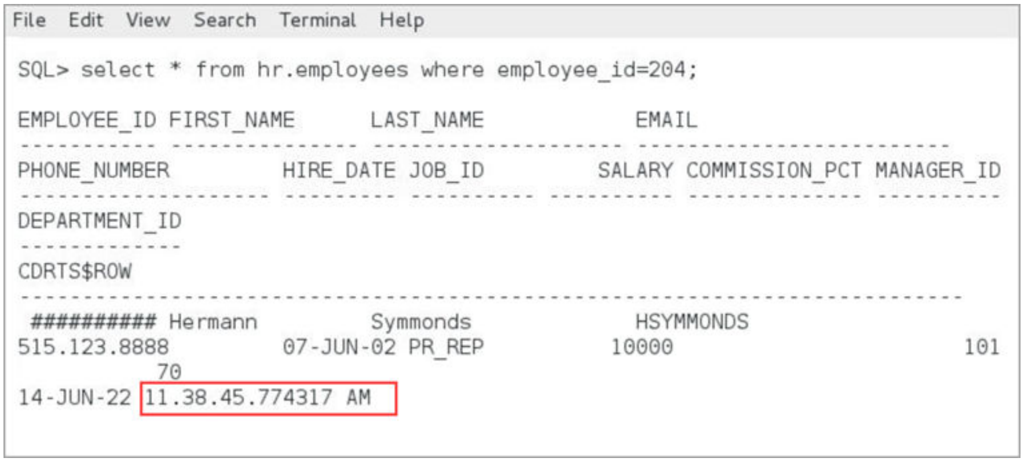

```
SQL> select * from hr.employees where employee_id=204;

EMPLOYEE_ID FIRST_NAME           LAST_NAME                 EMAIL
----------- -------------------- ------------------------- -------------------------
PHONE_NUMBER         HIRE_DATE JOB_ID         SALARY COMMISSION_PCT MANAGER_ID
-------------------- --------- ---------- ---------- -------------- ----------
DEPARTMENT_ID
-------------
CDRTS$ROW
---------------------------------------------------------------------------
 ########## Hermann              Symmonds                  HSYMMONDS
515.123.8888         07-JUN-02 PR_REP          10000                       101
           70
14-JUN-22 11.38.45.774317 AM
```

You can note the timestamp for this transaction: 11.38.45.774317 AM.

Now, let's check the timestamp on **DBWEST**:

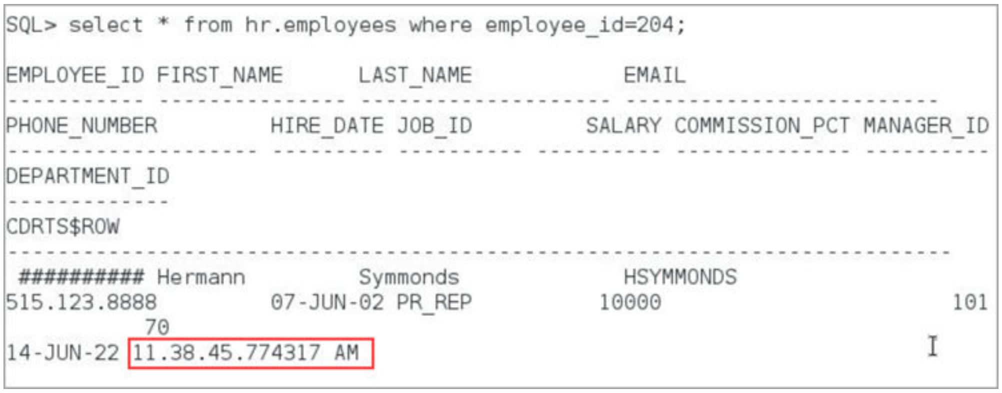

```
SQL> select * from hr.employees where employee_id=204;

EMPLOYEE_ID FIRST_NAME           LAST_NAME                 EMAIL
----------- -------------------- ------------------------- -------------------------
PHONE_NUMBER         HIRE_DATE JOB_ID         SALARY COMMISSION_PCT MANAGER_ID
-------------------- --------- ---------- ---------- -------------- ----------
DEPARTMENT_ID
-------------
CDRTS$ROW
---------------------------------------------------------------------------
 ########## Hermann              Symmonds                  HSYMMONDS
515.123.8888         07-JUN-02 PR_REP          10000                       101
           70
14-JUN-22 11.38.45.774317 AM
```

As the conflict is resolved, the timestamp shows the same data on both PDBs.

Run the following command to check if the conflicts were detected and resolved on the Oracle GoldenGate side. Here's the command to check this:

```
STATS REPLICAT repe, REPORTCDR
```

The output for this command displays the following:

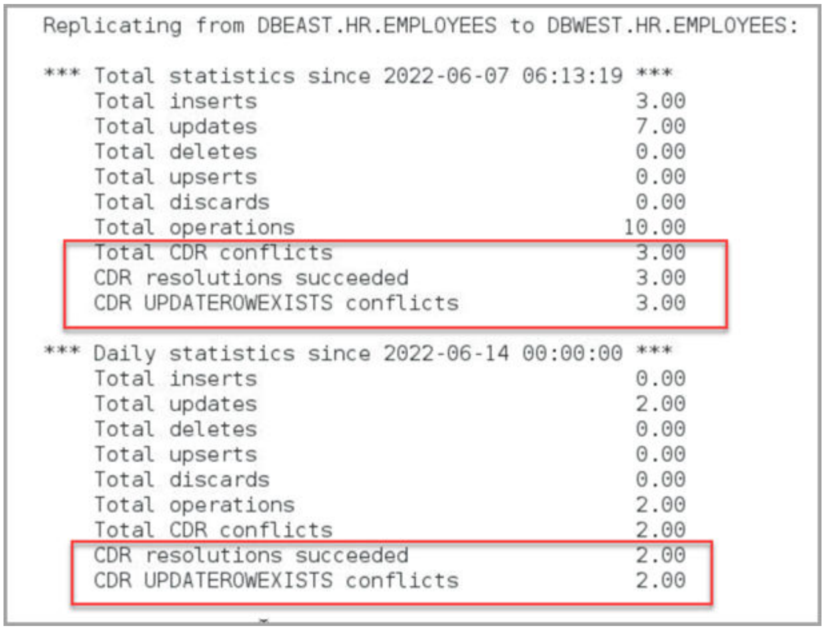

```
  Replicating from DBEAST.HR.EMPLOYEES to DBWEST.HR.EMPLOYEES:

  *** Total statistics since 2022-06-07 06:13:19 ***
      Total inserts                                3.00
      Total updates                                7.00
      Total deletes                                0.00
      Total upserts                                0.00
      Total discards                               0.00
      Total operations                            10.00
      Total CDR conflicts                          3.00
      CDR resolutions succeeded                    3.00
      CDR UPDATEROWEXISTS conflicts                3.00

  *** Daily statistics since 2022-06-14 00:00:00 ***
      Total inserts                                0.00
      Total updates                                2.00
      Total deletes                                0.00
      Total upserts                                0.00
      Total discards                               0.00
      Total operations                             2.00
      Total CDR conflicts                          2.00
      CDR resolutions succeeded                    2.00
      CDR UPDATEROWEXISTS conflicts                2.00
```

As shown in this statistical report, there were 3 conflicts and 2 of them were resolved. The `UPDATEROWEXISTS` conflict type is used for resolution.

Also see CDR Example 1: All Conflict Types with USEMAX, OVERWRITE, DISCARD in the *Oracle GoldenGate Administration Guide*.

You can also see this report from the web interface:

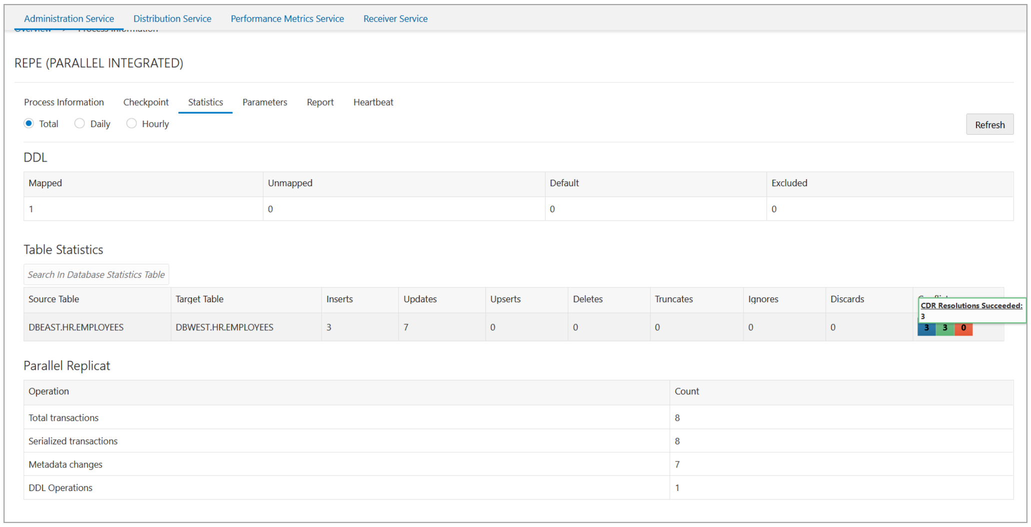

REPE (PARALLEL INTEGRATED)

DDL

Mapped	Unmapped	Default	Excluded
1	0	0	0

Table Statistics

Source Table	Target Table	Inserts	Updates	Upserts	Deletes	Truncates	Ignores	Discards	
DBEAST.HR.EMPLOYEES	DBWEST.HR.EMPLOYEES	3	7	0	0	0	0	0	3 3 0

CDR Resolutions Succeeded: 3

Parallel Replicat

Operation	Count
Total transactions	8
Serialized transactions	8
Metadata changes	7
DDL Operations	1

1.3 Switching from Nonintegrated Replicat to Parallel Nonintegrated Replicat

The process for switching to *parallel integrated* or parallel *nonintegrated Replicat* is the same for all Replicat modes. This topic describes the process to *switch from nonintegrated Replicat to parallel nonintegrated Replicat*.

Before Starting the Switching Process

1. Create a parallel nonintegrated Replicat process, **repe** that reads from the exisiting trail file:

   ```
   ADD REPLICAT repe, PARALLEL, EXTTRAIL ea, checkpointtable
   ggadmin.ggs_checkpoint1
   ```

   In this command, **repe** is the name of the Replicat. **ea** is the trail name. The trail name supplied while creating this Replicat is the same as the other Replicat in nonintegrated mode.

   > **Note:**
   >
   > If the checkpoint table is configured in GLOBALS, then there is no need to include the `checkpointtable` option with this command. If not, then use this option to provide the checkpoint table name.

2. Do ***not*** start the parallel nonintegrated Replicat (repe).
3. Stop the current nonintegrated Replicat, **repea**.

   ```
   STOP REPLICAT repea
   ```

4. On the target side, access the Replicat report file (.rpt) to know the values of the following components:

- Last applied CSN by the current nonintegrated Replicat process.
- Trail sequence and RBA of the exisitng Replicat process.

To access the details of the Replicat, run the command:

```
INFO REPLICAT repea DETAIL
```

The output for this command shows an output similar to the following:

```
Replicat   REPEA     Last Started 2022-06-16 04:21   Status STOPPED
Description          eastt
Checkpoint Lag       00:00:00 (updated 01:59:59 ago)
Log Read Checkpoint  File east/ea000000009
                     2022-06-14 04:38:34.084220  RBA 9382
Settings Profile     Default
Encryption Profile   LocalWallet

Current Log BSN value: (no data)

Last Committed Transaction CSN value: 50698907

  Extract Source                          Begin
End

  east/ea000000009                        2022-06-16 04:21  2022-06-14
04:38
  east/ea000000000                        * Initialized *   2022-06-16
04:21
  east/ea000000009                        2022-06-14 04:38  2022-06-14
04:38
  east/ea000000009                        2022-06-14 04:38  2022-06-14
04:38
  east/ea000000009                        2022-06-16 03:55  2022-06-14
04:38
  east/ea000000000                        * Initialized *   2022-06-16
03:55
  east/ea000000000                        * Initialized *   First
Record
  east/ea000000000                        * Initialized *   First
Record
  east/ea000000000                        * Initialized *   First
Record
  east/ea000000000                        * Initialized *   First
Record

Current directory    /scratch/preeshuk/ggtest/install_ogg21.3_210725/bin

Report file          /scratch/oggoradep/var/lib/report/REPEA.rpt
Parameter file       /scratch/oggoradep/etc/conf/ogg/REPEA.prm
Checkpoint file      /scratch/oggoradep/var/lib/checkpt/REPEA.cpr
Checkpoint table     DBEAST.GGADMIN.GGS_CHECKPOINT
Process file         /scratch/oggoradep/var/run/REPEA.pcr
Error log            /scratch/oggoradep/var/log/ggserr.log
```

Start the Switching Process

To start using the nonintegrated parallel Replicat, you need to alter it to port the content from the other Replicat. Use the following steps to perform this task:

1. Run the `ALTER REPLICAT` command as follows:

   ```
   ALTER REPLICAT replicat_name, EXTSEQNO extseqno,  EXTRBA extrba
   ```

   For example, for the Replicat repe, here's the command:

   ```
   ALTER REPLICAT repe, EXTSEQNO 9, EXTRBAm 9382
   ```

2. Start the newly created parallel nonintegrated Replicat process using the following command:

   ```
   START REPLICAT repe AFTERCSN csn_value
   ```

   For example:

   ```
   START REPLICAT repe AFTERCSN 50698907
   ```

This starts the Replicat at the specified CSN value in the trail file.

2 Configure Deployments

You can choose to set up a secure or non-secure deployment but whatever type you choose, all subsequent deployments of the same Service Manager must be of the same security type and cannot be changed afterwards.

A secure deployment makes RESTful API calls and transmits trail data between the Distribution Service and Receiver Service over SSL/TLS. You can choose to use an existing business cerificate from the Certificate Authority (CA) or create your own certificates.

For a non-secure deployment, the RESTful API calls occur over plain-text HTTP and conveyance between Distribution Service and Receiver Service is performed using the following protocols:

- Web Socket Protcol (`ws`)
- Secure Web Socket Protocol (`wss`)
- Oracle GoldenGate Classic Architecture Manager port (`ogg`)

For secure deployments, use the `wss` protocol.

This section describes the steps to configure deployments using Oracle GoldenGate Configuration Assistant (OGGCA) utility.

Oracle recommends using a secure configuration within Oracle GoldenGate MA.

See this Configuring a Secure Deployment with OGGCA that demonstrates setting up a secure deployment using OGGCA:

To set up additional configuration options for security, you can also see Configure Reverse Proxy with NGINX to Access Oracle GoldenGate Microservices.

Topics:

- Prepare the Database
- Install Oracle GoldenGate Microservices
- Add a Deployment
- Configure Reverse Proxy with NGINX to Access Oracle GoldenGate Microservices
- Remove a Deployment

2.1 Prepare the Database

Configure the database for Oracle GoldenGate replication.

To prepare your database for Oracle GoldenGate, ensure that your database meets the requirements as outlined in *Installing Oracle GoldenGate* , *Using Oracle GoldenGate for Oracle Database* and *Using Oracle GoldenGate for Heterogeneous Databases* guides.

2.2 Install Oracle GoldenGate Microservices

Learn about the tasks to perform for setting up and using Oracle GoldenGate microservices after you complete installing Oracle GoldenGate Microservice Architecture.

This guide assumes that you have already completed installing Oracle GoldenGate Microservices Architecture. See Installing the Microservices Architecture for Oracle GoldenGate in *Installing Oracle GoldenGate* .

Topics:

2.3 Add a Deployment

Before You Begin

Before starting the OGGCA wizard to add a deployment for the *first-time*, make sure that you have decided on the following aspects:

- *Service Manager will run as a system service or not*: Deciding to run the Service Manager as a system service (daemon) implies that it will be automatically started and stopped when the operating system is started and stopped. See Start and Stop Service Manager and Deployments.
- *Deployment would be added as a secure or non-secure deployment*: The security configuration options in OGGCA would require certificates and wallet configuration. See Single Deployment: Create Different Types of Certificates for a Secure Deployment.
- *Values for the Environment Variable*: Ensure that you are ready with the values for the environment variables to be configured for the database and operating system.
- *User authentication using Oracle GoldenGate local credentials or using Oracle Identity Cloud Service (IDCS)*: If you are using the local credentials for Oracle GoldenGate, then you can provide the username and password for logging into the Service Manager as the administrator. If using IDCS, then you need to have the credentials for logging in to the IDCS server.

Start OGGCA

After completing the Oracle GoldenGate MA installation, you can add one or more deployments using the Oracle GoldenGate Configuration Assistant (OGGCA) wizard. Adding a deployment is the first task in the process of setting up a data replication environment.

Deployments are managed from the Service Manager.

> **Note:**
>
> Oracle recommends that you have a *single Service Manager per host*, to upgrade the same Service Manager with new releases or patches and avoid redundant upgrade and maintenance tasks with new Oracle GoldenGate releases.

Use OGGCA to add deployments to a Service Manager. Oracle GoldenGate deployments across different servers are the *endpoints* for setting up the distribution path for data replication. To start the OGGCA wizard:

1. Navigate to the `$OGG_HOME/bin` directory.
2. Run the `oggca.sh` program on UNIX or `oggca.bat` on Windows.

Select Service Manager Options

1. Select the **Create a New Service Manager** option if you are running OGGCA for the *first time*. When you run OGGCA for the first time, the **Existing Service Manager** option is disabled. If it's not the first time, then you can choose the **Existing Service Manager** option, which would load the Service Manager port and other settings as configured for the existing Service Manager. The deployment would be added to this Service Manager. In most configurations, there is only one Service Manager to manage multiple deployments.
2. For a new Service Manager, browse and enter the directory that you want to use for your deployment in the **Service Manager Deployment Home** text box. Oracle recommends that you create a `ServiceManager` directory within the deployment sub-directory structure to store the Service Manager files.
3. Enter the connection details for the Service Manager:
   a. **Listening hostname/address**: Enter a hostname such as localhost or the IP address of the server where Service Manager will run.
   b. **Listening Port**: Enter a unique port number that the Service Manager will listen on, or choose the port already in use if selecting an existing Service Manager.
4. (Optional) Select the option **Register the Service Manager as a system service (daemon)** to avoid manually starting and stopping it if the machine is rebooted. If there is an existing Service Manager registered as a service and you select a new Service Manager to register as a service, an alert is displayed indicating that you cannot register the new one as a system service. All other Service Managers are started and stopped using scripts installed in the `bin` directory of the deployment.

   You cannot register an existing Service Manager as a system service. Enter a unique port number that the Service Manager will listen on, or choose the port already in use, if selecting an existing Service Manager.
5. (Optional) Select the **Integrate with XAG** option to integrate your deployment with an Oracle Grid Infrastructure for Oracle Database. This is only available for Oracle database in a cluster environment. This option cannot be used when running your Service Manager as a system service.
6. Click **Next**.

Configuration Options

In the **Configuration Options** step, you can add or remove deployments.

You can only add or remove one deployment for one Service Manager at a time.

> **Note:**
>
> Ensure that your Service Manager is up and running prior to launching OGGCA.

Deployment Details

1. Enter the deployment name using these conventions:
   - Must begin with a letter.
   - Can be a standard ASCII alphanumeric string not exceeding 32 characters.
   - Cannot include extended ASCII characters.
   - Special characters that are allowed include underscore (‘_’), forward slash (‘/’), dash (‘-’), period (‘.’).
   - Cannot be “ServiceManager”.
2. Select the **Enable FIPS** check box to enable Oracle GoldenGate services to use FIPS-compliant libraries.
3. (Oracle Database only) Select **Enable Sharding** to use the database sharding feature in your deployment. The schema must be `ggadmin`.
4. Enter or select the Oracle GoldenGate installation directory. If you have set the `$OGG_HOME` environment variable, the directory is automatically populated. Otherwise, the parent directory of the `oggca.sh` (Linux) or `oggca.bat` (Windows) script is used.
5. Click **Next**.

Select Deployment Directories

1. Enter or select a deployment directory where you want to store the deployment registry and configuration files. When you enter the deployment directory name, it is created if it doesn’t exist. Oracle recommends that you do *not* locate your deployment directory inside your `$OGG_HOME` and that you create a separate directory for easier upgrades. The additional fields are automatically populated based on the specified deployment directory.

   > **Note:**
   >
   > The deployment directory name (user deployment directory) needs to be different than the directory name chosen in the first screen (Service Manager deployment directory).

2. You can customize the deployment directories so that they are named and located differently from the default.
3. Enter or select different directories for the various deployment elements.
4. Click **Next**.

Environment Variables

Enter the requested values for the environment variables. Double-click in the field to edit it. You can copy and paste values in the environment variable fields. Make sure that you tab or click outside of the field after entering each value, otherwise it’s not saved. If you have set any of these environment variables, the directory is automatically populated.

OGG_HOME
The directory where you installed Oracle GoldenGate. This variable is fixed and cannot be changed.

> **Note:**
>
> On a Windows platform, ensure that there's no space in the `OGG_HOME` directory path otherwise OGGCA will not run.

IBMCLIDRIVER
Valid for DB2 z/OS.

Specifies the location where the IBM Data Server Driver for ODBC and CLI (IBMCLIDRIVER) software is installed.

LD_LIBRARY_PATH
This variable is used to specify the path to search for libraries on UNIX and Linux. It may have a different name on some operating systems, such as `LIBPATH` on IBM AIX on POWER Systems (64-Bit), and `SHLIB_PATH` on HP-UX. This path points to the Oracle GoldenGate installation directory and the underlying instant client directory by default.
If you are using User Exits, then append the `LD_LIBRARY_PATH` variable with the path to the additional shared libraries of the User Exit.

TNS_ADMIN
Valid for Oracle database.
This variable is recommended and points to the directory location containing `tnsnames.ora`, which has the database connection details. If this variable is not set, Oracle GoldenGate looks for `$HOME/.tnsnames.ora` or `/etc/tnsnames.ora` You need to create the `tnsnames.ora` file with the connection data and place it in the `$OGG_HOME/etc`. Here's a sample structure of the file:

```
# tnsnames.ora Network Configuration File:
# Generated by Oracle configuration tools.

LISTENER_ORCL19 =
  (ADDRESS = (PROTOCOL = TCP)(HOST = eastdb.us.oracle.com)(PORT = 1521))

ORCL =
  (DESCRIPTION =
   (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP)(HOST = eastdb.oracle.com)(PORT = 1521))
   )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = orcl.us.oracle.com)
    )
  )
```

For example: `TNS_ADMIN=/u01/app/oracle/network/admin`

STREAMS_POOL_SIZE
For Oracle Database Sharding only. This variable is mandatory for sharded databases. Use the default or set your pool size value that is at least 1200MB.

TZ
Valid for MySQL.
Use the following query to determine the timezone of the MySQL database:

```
SELECT @@global.time_zone;mysql> select
@@global.time_zone;+--------------------+| @@global.time_zone |
+--------------------+| SYSTEM |+--------------------+1 row in set
(0.00 sec)2
```

If it returns the timezone as `SYSTEM`, then it indicates that the database timezone is the same as the system timezone.
To know the system timezone, you can run the following query on your MySQL database:

```
mysql> select @@system_time_zone;+--------------------+|
@@system_time_zone |+--------------------+| UTC |+--------------------
+1 row in set (0.00 sec)3
```

Make sure that the database timezone and the timezone of the system where the Oracle GoldenGate instance is running are the same.
Alternatively, you can set the `TZ` variable for the deployment to the same value as the database timezeone. Use the following command to check the `TZ` variable on the shell where you are going to invoke `TZ`:

```
linux# echo $TZa
```

Use the following command to set `TZ` on the shell where you start your Oracle GoldenGate command line (Admin Client or GGSCI) from:

```
linux# export TZ UTC
```

ODBCINI
Valid for Oracle GoldenGate installed on Linux for PostgreSQL databases.
Specifies the full path of the ODBC file used to store Data Source Names (DSN) for connectivity to a PostgreSQL database. For example, `ODBCINI=/etc/odbc.ini`

JAVA_HOME
If this variable is present during deployment creation, it will automatically be populated, otherwise you can set it to be:

```
export JAVA_HOME=$OGG_HOME/jdk
```

You can add additional environment variables to customize your deployment or remove variables.

Click **Next**.

Administrator Account

To choose between Identity Cloud Service (IDCS) or local credential setup, define your Service Manager administrator user.

> **Note:**
>
> The option to set up IDCS-enabled administrator account is not applicable when you run OGGCA for the first time. Only after creating and enabling the Authorization Profile, you can set up the Administrator Account for accessing IDCS. See Enabling Authorization Profile.

1. Enter a user name and password that you want to use to sign in to the Oracle GoldenGate MA Service Manager and the other services. This user is the security user for this deployment.

   If you are adding a deployment to an exisiting Service Manager and intend to use IDCS (as your external Identity Provider) for user authentication, then specify the user credentials for the IDCS server. As a prerequisite to providing the credentials for accessing the IDCS server, you need to enable the Authorization Profile from the Service Manager deployment.

   > **Note:**
   >
   > For Administrator Account, you must enter a user and password for a provisioned external IDP identity that is mapped to the SECURITY group previously configured for the Service Manager deployment.

   Select the **Enable strong password policy in the new deployment** checkbox to ensure setting a highly secure password for your user account. This password policy applies for your localCredentialStore only but not for IDCS default settings. See Manage Oracle Identity Cloud Service Password Policies in *Administering Oracle Identity Cloud Service* guide.

   The strong password policy for **localCredentialStore** has the following requirements:

   - At least one lowercase character [a...z]
   - At least one upposercase character [A...Z]
   - At least one digit [0...9]
   - At least one special character [- ! @ % & * . #]
   - The length should be between 8 and 30 characters.

   For details on the different types of users, see How to Add Users. If you are using an existing Service Manager, you must enter the same log in credentials that were used when adding the first deployment.

2. Select the check box that allows you to enable a strong password policy for your new deployment. If you select this option, then the password must adhere to restrictions, otherwise an error occurs, which requires you to specify a stronger password.

3. Click **Next**.

Local Administrator Account Credentials

On this screen, enter the user credentials for the local administrator for the new deployment. If you want to enable IDCS for this new deployment, you can do so by enabling the authorization profile.

> **Note:**
>
> If Service Manager is enabled for IDCS, it can continue to manage the new deployment, which uses local administrator credentials, even if the new deployment is not enabled for IDCS.

Security Options

1. You can choose whether or not you want to secure your deployment. Oracle recommends that you enable SSL/TLS security.

   If you do not want to use security option on the source endpoint, deselect the check box.

2. When you deselect the SSL/TLS check box, the option **This non-secure deployment will be used to send trail data to a secure deployment** stays enabled. Select this check box to set up a secure target deployment to communicate with a non-secure source deployment. In this case, certificates are required for the client only.

   However, you must enable security if configuring for Oracle GoldenGate sharding support for Oracle Database.

3. For the Server (wallet or certificate), select the option to use a Wallet or Certificate. Provide the location of the wallet directory and if you are using an existing wallet, it must have the appropriate certificates already imported into it. If you choose to use a certificate, enter the corresponding pass phrase.

   When using a self-signed certificate, a new Oracle Wallet is created in the new deployment and these certificates are imported into it. For certificates, enter the location of the private key file and the pass phrase. The private key files must be in the `PKCS#8` format.

4. For the Client side, select either the wallet directory or certificate. Provide the wallet directory on the client side or the certificate details for the client. If you select the **This non-secure deployment will be used to send trail data to a secure deployment**, then you only need to specify the client side details (wallet or certificate of the target deployment). This option is useful when the Distribution Service from the source deployment is unsecured whereas the Receiver Service on the target deployment is secured. So, the sender may be configured for public access while the Receiver Service requires authentication and authorization, which is established using PKI before the incoming data is applied.

   For more information, see Single Deployment: Create Different Types of Certificates for a Secure Deployment.

   Also see: #unique_28.

5. Click **Next**.

Advanced Security Settings

(If Security is enabled) On the page, the encryption options TLS 1.1 and TLS 1.2 are available. TLS 1.2 is selected by default.

When you open the Advanced Security Settings for the first time with TLS 1.2, the following cipher suites are listed:

```
TLS_RSA_WITH_AES_128_CBC_SHA256
TLS_RSA_WITH_AES_128_GCM_SHA256
TLS_RSA_WITH_AES_256_CBC_SHA256
TLS_RSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384
TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384
TLS_ECDH_ECDSA_WITH_AES_128_CBC_SHA256
TLS_ECDH_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDH_ECDSA_WITH_AES_256_CBC_SHA384
TLS_ECDH_ECDSA_WITH_AES_256_GCM_SHA384
TLS_ECDH_RSA_WITH_AES_128_CBC_SHA256
TLS_ECDH_RSA_WITH_AES_128_GCM_SHA256
TLS_ECDH_RSA_WITH_AES_256_CBC_SHA384
TLS_ECDH_RSA_WITH_AES_256_GCM_SHA384
```

1. Use the arrows to add or remove cipher suites or the **Up** and **Down** to reorder how the cipher suites will be applied.
2. Click **Next**.

Sharding Options

If Sharding was enabled in the previous step, then you can configure the sharding options on this screen.

1. Locate and import your Oracle GoldenGate Sharding Certificate. Enter the distinguished name from the certificate that will be used by the database sharding code to identify itself when making REST API calls to the Oracle GoldenGate MA services.
2. Enter a unique name for the certificate.
3. Click **Next**.

Port Settings

1. Enter the Administration Service port number, and then when you leave the field the other port numbers are populated in ascending numbers. Optionally, you can enter unique ports for each of the services.

2. Select **Enable Monitoring** to use the Performance Metrics Service.
3. Click inside the Performance Metrics Service port fields to populate or enter the ports you want to use. Ensure that you choose available ports for TCP.

   Select the UDP port for performance monitoring. The option to select the UDP port is displayed only with deployments on Windows and other operating systems that don't support UDS communication with Performance Metric Service. See Supported Operating Systems for UDS.

   You can change the TCP port from the Service Manager console after the deployment is done. For more information on `PMSRVR`, see `ENABLEMONITORING`.
4. Select the type of datastore that you want the Performance Metrics Service to use, the default Berkeley Database (BDB) data store or Open LDAP Lightning Memory-Mapped Database (LMDB). You can also designate the Performance Monitor as a Critical Service if integrating the Service Manager with XAG.

   For BDB informtion, see Oracle Berkeley DB 12*c* Release 1. For LMDB information, see http://www.lmdb.tech/doc/.
5. Select the location of your datastore. BDB and LMDB are in-memory and disk-resident databases. The Performance Metrics Service uses the datastore to store all performance metrics information.
6. Click **Next**.

> **Note:**
>
> The `oggca` utility validates whether or not the port you entered is currently in use or not.

Replication Settings

1. Enter the Oracle GoldenGate default schema name such as `ggschema` that you want to use to store the replication objects such as the checkpoint and heartbeat tables.

   > **Note:**
   >
   > OGGCA does not connect to the database, so it cannot validate the schema. The schema specified in OGGCA is written to the GLOBALS file as a default schema. When creating an Extract, if you do not specify a replication schema, Extract will use that default schema.
2. Click **Next**.

Summary

1. Review the detailed configuration settings of the deployment before you continue.
2. (Optional) You can save the configuration information to a response file. Oracle recommends that you save the response file. You can run the installer from the command line using this file as an input to duplicate the results of a successful configuration on other systems. You can edit this file or a new one from the provided template.

> **Note:**
>
> When saving to a response file, the administrator password is not saved for security reasons. You must edit the response file and enter the password if you want to reuse the response file for use on other systems.

3. Click **Finish** after reviewing the response file.
4. Click **Next**.

Configure Deployment

Displays the progress of the deployment creation and configuration.

1. If the Service Manager is being registered as a service, a pop-up appears that directs you how to run the script to register the service. OGGCA verifies that these scripts have been run. If you did not run them, you are queried if you want to continue. When you click **Yes**, the configuration completes successfully. When you click **No**, a temporary failed status is set and you click **Retry** to run the scripts.
2. Click **OK** after you run the script to continue and then click **Next**.

Finish

Click **Close** to exit OGGCA.

Topics:

- Single Deployment: Create Different Types of Certificates for a Secure Deployment
- Two Deployments: Create External, Trusted Server and Client Certificates

2.3.1 Single Deployment: Create Different Types of Certificates for a Secure Deployment

These certificates are used if you have one deployment having a Distribution path from the Distibution and Reciver Service within a single (or same) deployment.

Here's how you can create client and server certificates to set up a secure Oracle GoldenGate Microservices Architecture deployment.

Topics:

- Create a Self-Signed Trusted (Root) Certificate
- Create Server Certificates
- Create a Client Certificate
- Set Up Trusted Certificates

2.3.1.1 Create a Self-Signed Trusted (Root) Certificate

You may apply your existing trusted certificate or use the `orapki` in the `OGG_HOME/bin` directory.

Here's an example of how you can create a root certificate using `orapki`:

1. Create a directory to store your wallets and certificates. For example, `~/` *`wallet_directory`*.
2. Create an automatic login wallet. This example uses `root_ca` for the wallet name.

   ```
   orapki wallet create -wallet ~/wallet_directory/root_ca -auto_login
   -pwd welcome123
   ```

3. In the `orapki` command to create self-signed (root user) certificate, specify the-`sign_alg sha256` option.
4. In `orapki` wallet:

   ```
   orapki wallet add -wallet ~/wallet_directory/root_ca -dn
   "CN=RootCA"-addext_basic_cons -pathlen 10 -keysize 2048 -
   self_signed -validity 7300 -pwd welcome123 -sign_alg sha256
   ```

5. Export the certificate to a `.pem` file.

   ```
   orapki wallet export -wallet ~/wallet_directory/root_ca -dn
   "CN=RootCA" -cert ~/wallet_directory/rootCA_Cert.pem -pwd welcome123
   ```

The wallet creation is complete.

2.3.1.2 Create Server Certificates

The following steps are an example of how you can create a sever certificate using a root certificate named `root_ca`.

1. Create a directory to store your wallets and certificates. For example, `~/` *`wallet_directory`*.
2. Create an automatic login server wallet.

   ```
   orapki wallet create -wallet ~/wallet_directory/$(hostname) -
   auto_login -pwd welcome123
   ```

   Enter the password for the server when prompted.
3. Add a Certificate Signing Request (CSR) to the server’s wallet.

   ```
   orapki wallet add -wallet ~/wallet_directory/$(hostname) -dn "CN=$
   (hostname)" -addext_basic_cons -pathlen 10 -keysize 2048 -pwd
   welcome123
   ```

> **Note:**
>
> The `addext_basic_cons -pathlen 10` option is important as it is used to apply the later certificate into another secure store, when setting certificates for two different deployments. See Two Deployments: Create External, Trusted Server and Client Certificates.

4. Export the CSR to a `.pem` file.

   ```
   orapki wallet export -wallet ~/wallet_directory/$(hostname) -dn "CN=$
   (hostname)" -request ~/wallet_directory/servername_req.pem -pwd welcome123
   ```

5. Using the CSR, create a signed server or client certificate and sign it using the root certificate. Assign a unique serial number to each certificate.

   ```
   orapki cert create -wallet ~/wallet_directory/root_ca -request ~/
   wallet_directory/servername_req.pem -cert ~/wallet_directory/
   servername_Cert.pem -serial_num 20 -validity 375  -sign_alg sha256
   ```

6. Add the root certificate into the client's or server's wallet as a trusted certificate.

   ```
   orapki wallet add -wallet ~/wallet_directory/$(hostname) -trusted_cert -
   cert ~/wallet_directory/rootCA_Cert.pem -pwd welcome123
   ```

7. Add the server or client certificate as a user certificate into the client's or server's wallet.

   ```
   orapki wallet add -wallet ~/wallet_directory/$(hostname) -user_cert -cert
   ~/wallet_directory/servername_Cert.pem  -pwd welcome123
   ```

The wallet creation is complete.

2.3.1.3 Create a Client Certificate

The following steps are an example of how you can create a Distribution Service user certificate:

1. Create a directory to store your wallets and certificates. For example, `~/wallet_directory`.
2. Create an automatic login client wallet. This example uses `dist_client` for the wallet name.

   ```
   orapki wallet create -wallet ~/wallet_directory/dist_client -auto_login -
   pwd welcome123
   ```

3. Add a CSR to the wallet.

   ```
   orapki wallet add -wallet ~/wallet_directory/dist_client -dn
   "CN=dist_client" -keysize 2048 -pwd welcome123
   ```

4. Export the CSR to a `.pem` file.

   ```
   orapki wallet export -wallet ~/wallet_directory/dist_client -dn
   "CN=dist_client" -request ~/wallet_directory/dist_client_req.pem -
   pwd welcome123
   ```

5. Using CSR, create a signed server or client certificate and sign it using the root certificate. Assign a unique serial number to each certificate.

   ```
   orapki cert create -wallet ~/wallet_directory/root_ca -request ~/
   wallet_directory/dist_client_req.pem -cert ~/wallet_directory/
   dist_client_Cert.pem -serial_num 30 -validity 375 -pwd welcome123
   ```

6. Add the root certificate as a trusted certificate into the client’s or server’s wallet.

   ```
   orapki wallet add -wallet ~/wallet_directory/dist_client -
   trusted_cert -cert ~/wallet_directory/rootCA_Cert.pem -pwd
   welcome123
   ```

7. Add the server or client certificate as a user certificate into the client’s or server’s wallet.

   ```
   orapki wallet add -wallet ~/wallet_directory/dist_client -user_cert
   -cert ~/wallet_directory/dist_client_Cert.pem -pwd welcome123
   ```

The wallet creation is complete.

2.3.1.4 Set Up Trusted Certificates

There are two types of TLS connections. To use TLS, there are certain requirement for the certificate trust chain.

The `wss` communication protocol is used in the Distribution Service for the Distribution Path to meet the needs of secure communication using TLS in Oracle GoldenGate Microservices Architecture.

Setting up the server's CA certificate as a Trusted Certificate for External Identity Provider

To work with an external Identity Provider (IDP) such as IDCS, you need to upload the IDP server's (IDCS) CA certificate as a trusted certificate. See Add and Manage Certificates for a Deployment.

Distribution Service and Receiver Service

Both the Distribution Service and Receiver Service need certificates. The Distribution Service uses the certificate in the client wallet location under outbound section. The location of that wallet can be found in the `deploymentConfiguration.dat` file under *`deployment_home`*`/etc/conf`.

The certificates in both wallets need to be trusted by each other, so either both need to have commercial certificates issued by Classic Architecture, or they have to trust each other for self-signed certificates.

For self-signed certificates, you can choose from one of the following:

- Have both certificates signed by the same root certificate. (`rootCA`)
- The other side's certificate is added to the local wallet as a trusted certificate

For the Receiver Service, the certificate is in the wallet for the local wallet location, which is also in the `deploymentConfiguration.dat` file.

On the Distribution Service, if the hostname used in the Receiver Service's certificate can't be routed correctly, `/etc/hosts` file should be updated with the correct IP address for that host. The Distribution Service will use this IP address to communicate with the Receiver Service once it accepts the certificate from the Receiver Service.

Using the Reverse Proxy (Nginx) with the Distribution Service and Receiver Service

You only need to add the Nginx certificate to the Distribution Service's client wallet as a trusted certificate. Usually the certificate used by Nginx is self-signed. If it is issued by Classic Architecture, then there is no need to perform this step.

The host name in the Nginx certificate should also be routable. If not, on the Distribution Service, `/etc/hosts` file needs to be updated to reflect the correct IP address for that host name.The Distribution Service will use the host name in the certificate to communicate to the target. If the Nginx certificate doesn't have a valid host name in it, but has a Subject Alternative Name record, then the host name is the DNS name there.

2.3.2 Two Deployments: Create External, Trusted Server and Client Certificates

Each system (deployment) has its own set of Root, server, and client certificates, which are created using the `orapki` utility. These certificates can be part of wallets such as wallet_A and wallet_B.

In addition to these certificates, there is another set of external (extern) certificates for situations where the distribution path needs to be established between different source and target deployments, such as source A and target B.

Here's how you can create trusted, server, and client certificates for two different secure Oracle GoldenGate Microservices Architecture deployments:

- Add a Target Server Certificate as a CA Certificate

2.3.2.1 Add a Target Server Certificate as a CA Certificate

Use the following steps to create and manage server certificates as CA certificates for a target deployment that is different from the source deployment.

From these steps, you will be able to create a client certificate **client_src_to_trg** of the type **rootCA_extern**, which are generated using OpenSSL.

Source Deployment
At the source deployment side, perform the following tasks:

1. In the Service Manager, navigate to the Certificate Management page.
2. Under the **CA Certificates Shared** section, click the plus sign (+) to add the certificate of the target server. This is the server certificate from the target side that was earlier

created with orapki for the initial setup. See Add and Manage Certificates for a Deployment for steps to add the certificate.

> **Note:**
>
> The source Distribution Service must trust the server certificate used by the target. This needs to be added to the source secure store.

3. Under the **Client Certificates** section, click the plus sign (+) to add the client certificate (**client_SRC_to_TRG**), which will be used for the distribution path between source and target side.

   You might notice that this client certificate is signed by another trusted Root certificate (**rootCA_extern**). It is a client certificate that is created outside of the initial setup that was created earlier with orapki.

   Both, the root- and client certificates that are signed by this Root, are independent of the certificates from the initial deployment of the Oracle GoldenGate source and target instances.

Here's a sample of the client certificate configuration (client_src_to_trg.cfg) file:

```
[ req ]
default_bits = 4096
default_md = sha512
prompt = no
encrypt_key = no
distinguished_name = req_distinguished_name
[ req_distinguished_name ]
commonName = "client_src_to_trg"
[ my_extensions ]
```

Here are sample rootCA_extern.cfg configuration file:

```
[ req ]
default_bits = 4096
default_md = sha512
prompt = no
encrypt_key = no
distinguished_name = req_distinguished_name
req_extensions = v3_req
x509_extensions = v3_ca
x509_extensions = usr_cert
[ req_distinguished_name ]
#countryName = "US"
#stateOrProvinceName = "CA"
#localityName = "Redwood City"
#streetAddress = "400 Oracle Pkwy"
#organizationName = "Oracle USA Inc"
#organizationalUnitName = "Security"
commonName = "rootCA_extern"
#emailAddress = "rootsecurity@oracle.com"
```

```
[ v3_req ]
basicConstraints=CA:TRUE
[ v3_ca ]
basicConstraints=CA:TRUE
[ usr_cert ]
basicConstraints=CA:TRUE
[ my_extensions ]
```

Target Deployment

At the target deployment, perform the following tasks:

1. Under the CA Certificates section, click the plus sign (+) to add the trusted Root certificate (**rootCA_extern**) that was used to sign the previously added client certificate from the source side.

   > **Note:**
   >
   > The target (Receiver Service) must trust either the client certificate or the issuer of the client certificate. Therefore, it needs to be added to the target secure store.

2. From the Administration Service, add the user and role of the client that is used later for adding a distribution path. This user uses certificates for authorization. See Add Oracle GoldenGate Deployment Users from the Administration Service for steps to add users and roles.

Both the client certificate and the trusted Root certificate are independent from the certificates that were added in the initial deployment of the Oracle GoldenGate source and target instances. The certificates are created with OpenSSL commands.
Here's a sample rootCA certificate:

```
# rootCA certificate
openssl req -x509 -newkey rsa:4096 -keyout rootCA_extern.key -out
rootCA_extern.cert -days 73000 -nodes -config rootCA_extern.cfg

# client certificate
openssl req -new -newkey rsa:2048 -nodes -keyout client.key -out client.csr -
config client.cfg
openssl x509 -req -days 73000 -in client.csr -CA rootCA_extern.cert -CAkey
rootCA_extern.key -CAcreateserial -out client.cert
```

Creating the Distribution Path

After completing the setup of your **rootCA_extern** certificate on the target deployment, you can add a distribution path at the *source deployment* using the client certificate that was created for routing data from the source to the target system.
At the target deployment, you have to add a user with a specific role. This user is `CN=client_src_to_trg`.
SeeAdd a Distribution Path for steps to create your distribution path.

Note:

You will need to select the Target Authentication Method as Certificate to set up the distribution path between the source and target deployment.

2.4 Configure Reverse Proxy with NGINX to Access Oracle GoldenGate Microservices

Learn how to configure reverse proxy service using NGINX for accessing Oracle GoldenGate Microservices without using port numbers.

Reverse proxy enables accessing microservices using one single port (443) in a deployment. This enables encapsulation of the URL for microservices over an unsecure deployment.

Note:

Reverse proxy is optional, however, Oracle recommends that you ensure easy access to microservices and provide enhanced security.

You can run microservices in an unsecure deployment on loopback address and front it with an HTTP reverse proxy using the NGINX installation.

When sending trail files from Oracle GoldenGate Classic to Microservices environment that is configured with a reverse proxy, use a pump Extract from Oracle GoldenGate Classic with `SOCKSPROXY` option. When sending trail files from Oracle GoldenGate Microservices to Classic Architecture use the `ogg` protocol in the Distribution Service configuration.

See Connecting Classic to MA and Connecting MA to Classic in *Administering Oracle GoldenGate*.

Reverse Proxy Support

You can configure Oracle GoldenGate Microservices Architecture to use a reverse proxy. Oracle GoldenGate MA includes a script called `ReverseProxySettings` that generates configuration file for only the NGINX reverse proxy server.

For example, the Administration Service is available on `http://goldengate.example.com:9001` and the Distribution Service is on `http://goldengate.example.com:9002`. With reverse proxy, each of the microservices can simply be accessed from the single address. For example, `http://goldengate.example.com/distsrvr` for the Distribution Service. The URL is different for each service and is by name instead of by port.

You can use these options by running the `ReverseProxySettings` utility. Here are the options available with this utility:

-o or **--output**
The output file name. The default file name is ogg.conf.

-P or **--password**
A password for a Service Manager account.

-l or **--log**
Log file name and initiates logging. The default is no logging.

--trailOnly
Configure only for inbound trail data.

-t or **--type**
The proxy server type. The default is Nginx.

-s or **--no-ssl**
Configure without SSL.

-h or **--host**
The virtual host name for reverse proxy.

-p or **--port**
The reverse proxy port number. The defaults are 80 or 443.

-? or **--help**
Display usage information.

-u or **--user**
Name of the Service Manager account to use.

-v or **--version**
Displays the version.

These values are used when connecting to the Service Manager and are required when authentication is enabled.

You can use any reverse proxy service with MA. The following example provides a process that you can follow to configure other reverse proxy services in conjunction with the documentation for your proxy server.

Prerequisites

The following prerequisites provide details on the minimum requirements to configure an NGINX Reverse Proxy. Similar requirements may be required for your environment and reverse proxy if not using NGINX. Consult the documentation for your reverse proxy.

1. Install NGINX, see Install the NGINX Web Server and Proxy on Oracle Linux. For Oracle Linux, the command to install NGINX is:

   ```
   yum —y install NGINX
   ```

2. Check the JRE version to be JRE 8 or higher.
3. Install Oracle GoldenGate MA.
4. Create one or more active MA deployments.
5. Ensure that the Oracle user has sudo permissions.
6. Configure the PATH environment variable to include the NGINX installation directory path.

Configuring NGINX Reverse Proxy

An Oracle GoldenGate MA installation includes the `ReverseProxySettings` utility. The `ReverseProxySettings` utility is located in the *$OGG_HOME*`/lib/utl/reverseproxy` directory. To identify additional commands that can be used with the `ReverseProxySettings` utility, run the utility with the `--help` option:

```
$OGG_HOME/lib/utl/reverseproxy/ReverseProxySettings --help
```

To add the NGINX certificate to the Distribution Service's client wallet as a trusted certificate, see Set Up Trusted Certificates.

1. To generate a configuration file for NGINX Reverse Proxy, navigate to the location of the ReverseProxySettings utility:

   ```
   cd $OGG_HOME/lib/utl/reverseproxy
   ```

2. Run the ReverseProxySetting utility:

   ```
   ReverseProxySettings -u adminuser -P adminpwd -o ogg.conf http://
   localhost:9100
   ```

   In this code snippet, *`adminuser`* is the deployment user name and *`adminpwd`* is the deployment user password used to login to the deployment.

3. Replace the existing NGINX configuration with the configuration that was generated using the `ReverseProxySetting` utility for your MA deployment:

   ```
   sudo mv ogg.conf /etc/nginx/conf.d/nginx.conf
   ```

   However, this NGINX configuration isn't complete without the `events` section, and enclosing the `map` and `server` sections in `http`.

   Optionally, you can use the default `nginx.conf` file and add the generated `ogg.conf` by adding an `include` statement similar to this:

   ```
   include /etc/nginx/conf.d/ogg.conf;
   ```

   In this case, you must comment out the other `servers` section.

4. Generate a self-signed certificate for NGINX:

   ```
   sudo sh /etc/ssl/certs/make-dummy-cert /etc/nginx/ogg.pem
   ```

   For distribution paths to go through the reverse proxy, you need to use a valid certificate. It's better to specify the same certificate that the deployment is using to process incoming requests, otherwise, starting the path will fail with the next error in Distribution Service:

   ```
   2019-03-26T11:26:00.324-0700 ERROR| ERROR   OGG-10351  Oracle
   GoldenGate Distribution
    Service for Oracle:  Generic error -1 noticed. Error description -
   Certificate validation
   ```

```
 error: Unacceptable certificate from test00abc: application verification
failure. (A4)
```

5. Validate the NGINX configuration:

   ```
   sudo NGINX -t
   ```

   The output would show the following, if the command is successful:

   ```
   NGINX: the configuration file /etc/NGINX/NGINX.conf syntax is ok
   NGINX: configuration file /etc/NGINX/NGINX.conf test is successful
   ```

6. Reload NGINX with the new configuration:

   ```
   sudo NGINX -s reload
   ```

   If the changes for the configuration file are not loaded, stop and restart the proxy.

7. To test if you can access the microservices after NGINX is set up successfully, open the web browser.

8. Enter the proxy URL for the Service Manager using port number 443, similar to the following:

   http://dc.example.com:443

   This would open the Service Manager login page, from where you can access the other microservices also. If you want to directly access a microservice, you can enter the proxy URL for that microservice, as given in the `ogg.conf` file, generated previously.

Also see this video on configuring the NGINX reverse proxy.

SSL Termination

When there is an unsecure connection between the reverse proxy, which uses a TLS-based connection, and the origin server, it is referred to as reverse proxy SSL-termination.

> **Note:**
>
> In SSL-Termination the connections between the reverse proxy and the origin servers are unsecure.

However, SSL-bridging is also supported where the connections between the client and reverse proxy is secured and the connection between the reverse proxy and the origin server is also secured.

2.5 Remove a Deployment

You can remove a deployment using OGGCA or in silent mode.

Topics:

- Remove a Deployment: GUI
- Remove a Deployment: Silent Mode

2.5.1 Remove a Deployment: GUI

You can remove a deployment using the Oracle GoldenGate Configuration Assistnat wizard.

To remove a deployment:

> **Note:**
>
> When you remove a deployment or uninstall Oracle GoldenGate MA, the system does not automatically stop processes. As a result, you may have to stop processes associated with the deployment and you must clean files manually.

1. Run the Oracle GoldenGate Configuration Assistant wizard:

   ```
   $OGG_HOME/bin
   ```

2. Select **Existing Service Manager** from the **Select Service Manager Options** screen. Click **Next**
3. Select **Remove Existing Oracle GoldenGate Deployment** from the Configuration Options screen.
4. Select the deployment you need to remove from the **Deployment Name** list box. Also select the **Delete Deployment Files from Disk** check box if you want to remove all the deployment files (including configuration files) from the host.
5. Enter the Administration account user name and password and click **Next**.
6. See the list of settings that are deleted with the deployment and click **Finish**.

To remove a Service Manager:

1. Run Oracle GoldenGate Configuration Assistant wizard:

   ```
   $OGG_HOME/bin
   ```

2. Select **Existing Service Manager** from the **Select Service Manager Options** screen. Click **Next**.
3. If there are no other deployments to remove, then the option to remove the Service Manager is available in the drop down. Select **Remove Service Manager Deployment** from the Configuration Options screen.
4. Click **Finish**.

Files to be Removed Manually After Removing Deployment

It’s mandatory to delete some files manually only in case there's a Service Manager registered but you have to unregister it and register a new one. To remove files manually, you must have `root` or `sudo` privileges. The files to be deleted include:

Operating System	Files to be Removed Manually to Unregister an Existing Service Manager
Linux 6 **Note:** Linux 6 is not certified for Oracle GoldenGate 21c (21.3.0). This information may be required when trying to perform upgrades or downgrades.	• `/etc/init.d/OracleGoldenGate` • `/etc/rc.d/*OracleGoldenGate` • `/etc/rc*.d/*OracleGoldenGate` • `/etc/oggInst.loc`
Linux 7 and Linux 8	`/etc/systemd/system/OracleGoldenGate.service`

The following commands are executed to stop the Service Manager:

```
systemctl stop OracleGoldenGate
systemctl disable OracleGoldenGate *
```

Note:

If the Service Manager is not registered as a service (with or without the integration with XAG), OGGCA stops the Service Manager deployment, otherwise, a script called `unregisterServiceManager` is created, and when executed by the user, it runs the `systemctl` commands and deletes the mentioned files.

2.5.2 Remove a Deployment: Silent Mode

You can remove a deployment silently using the Oracle GoldenGate Configuration Assistant (oggca) from the Oracle GoldenGate Home bin directory.

By removing a deployment, you can delete various components of the deployment, including, Extracts, Replicats, paths, and configuration files. However, the Service Manager is not deleted.

To remove a deployment silently:

Note:

If the Service Manager is registered as a system service, removing a deployment silently will not unregister the service.

1. Ensure that you have a deployment response file. To get the deployment response file, run the OGGCA and the save the response file.

2. Update the following lines within the deployment response file:

   ```
   CONFIGURATION_OPTION=REMOVE
   ADMINISTRATOR_PASSWORD=********
   CREATE_NEW_SERVICEMANAGER=false
   DEPLOYMENT_NAME=deployment_name
   REMOVE_DEPLOYMENT_FROM_DISK=true
   ```

   In case of multiple deployments, you must specify the deployment name using the `DEPLOYMENT_NAME` field. You can use the `REMOVE_DEPLOYMENT_FROM_DISK` option to remove physical files and folders associated with deployment.

3. Run the OGGCA program from the following location using the `-silent` and `-responseFile` options. Providing the exact path to the deployment response is needed.

   ```
   $OGG_HOME/bin/oggca.sh -silent -responseFile
   path_to_response_file/response_file.rsp
   ```

   Example:

   ```
   $OGG_HOME/bin/oggca.sh -silent -responseFile
           /home/oracle/software/ogg_deployment.rsp
   ```

3

Manage Deployments from the Service Manager

Access the Service Manager to add deployments or edit existing ones.

The Service Manager is the primary watchdog service within Oracle GoldenGate MA that enables controlling the deployments and associated services running on the host machine.

The Service Manager can be configured in three different modes:

- Manually
- As a Daemon
- Integrated with XAG agent

For more information on setting up the Service Manager as a daemon service, see How to Create Secure and Non-Secure Deployments.

To start using your Oracle GoldenGate MA deployment, you need to connect to the Service Manager, initially:

1. Open a web browser and connect to the Service Manager that you created with Oracle GoldenGate Configuration Assistant (OGGCA).

   The URL is similar to `http://`*`host`*`:`*`port`*, where *`host`* is the name of the service or IP of the service that is running the Service Manager and *`port`* is the port number of the Service Manager. For a secure deployment, the URL is similar to `https://localhost:9001`.

   You can also connect to the Service Manager, using the Admin Client as shown in the following example:

   ```
   OGG Not Connected>Connect https://east.oracle.com:9000 as deployment ogg
   password admin
   OGG Connected>
   ```

   See the `CONNECT` command for details.

2. Enter the user name and password you created during deployment and sign in.

From the Service Manager, you can perform the following tasks:

- Check if the deployments and the corresponding microservices are up and running.
- Start, stop, and restart deployments, microservices, and the Service Manager
- Configure deployment details, environment variables
- Create and manage certificates for a deployment
- Add Oracle GoldenGate users with different roles

Apart from these, there are other functions that you can perform from the Service Manager.

Topics:

- Quick Tour of the Service Manager Overview Page
- Enable and Use Debug Logging
- Read the Log Information
- Add and Manage Certificates for a Deployment
- Add Users from the Service Manager

3.1 Quick Tour of the Service Manager Overview Page

The Service Manager allows you to perform key operations on your deployment such as check the status of all microservices, start and stop them, create users, manage certificates.

Open a web browser and enter the Service Manager URL. Login with the administrator account user credentials you specified while setting up the deployment in OGGCA.

The Service Manager home page is a dashboard where you can see the microservices and deployments listed in the **Services** and **Deployment** sections.

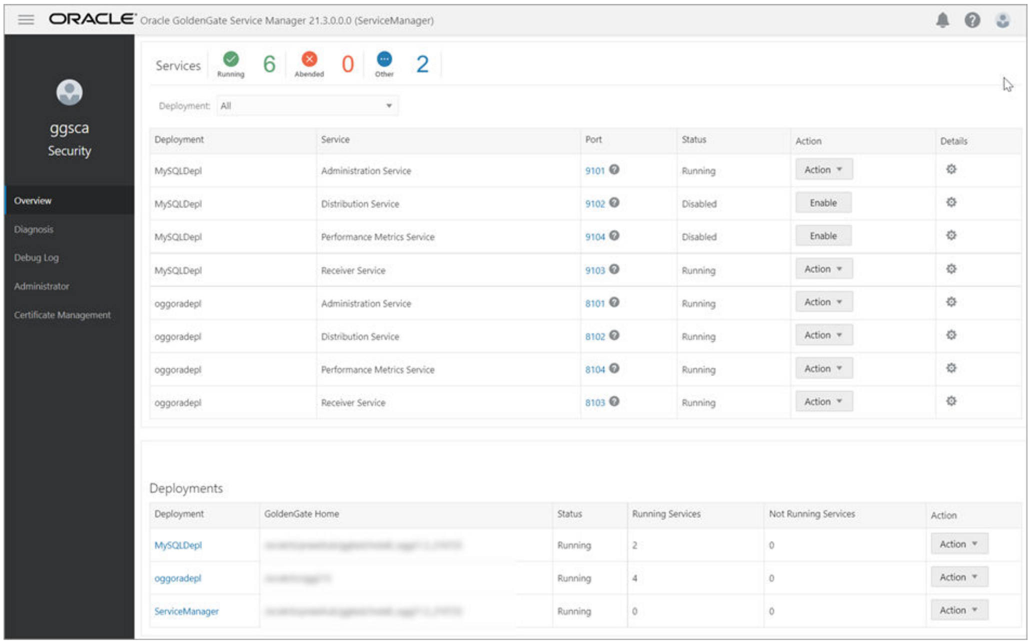

Deployment	Service	Port	Status	Action	Details
MySQLDepl	Administration Service	9101	Running	Action	⚙
MySQLDepl	Distribution Service	9102	Disabled	Enable	⚙
MySQLDepl	Performance Metrics Service	9104	Disabled	Enable	⚙
MySQLDepl	Receiver Service	9103	Running	Action	⚙
oggoradepl	Administration Service	8101	Running	Action	⚙
oggoradepl	Distribution Service	8102	Running	Action	⚙
oggoradepl	Performance Metrics Service	8104	Running	Action	⚙
oggoradepl	Receiver Service	8103	Running	Action	⚙

Deployments

Deployment	GoldenGate Home	Status	Running Services	Not Running Services	Action
MySQLDepl	[redacted]	Running	2	0	Action
oggoradepl	[redacted]	Running	4	0	Action
ServiceManager	[redacted]	Running	0	0	Action

From the **Services** section, you can:

- Access the microservices using the links provided in the Port column.
- Check the state of the microservices.
- Start/stop, disable/enable microservices from the Actions column.
- View and edit the details include the service configuration and restart options from the Details column. See View and Edit Services Configuration.

From the **Deployments** section:

- View the Oracle GoldenGate home location from the GoldenGate home column.
- Check the state of the deployment from the Status column.

- Start, stop, and restart a deployment from the Action column. See Start and Stop Service Manager and Deployments.
- View and edit the following details for a deployment by clicking the deployment name:
  - View and edit the deployment details.
  - Configure environment variables.
  - Add and manage certificates associated with the deployment.
  - Add and manage authorization profiles. By default, the localCredentialStore profile is added.

  See Modify Deployment Details and Configuration.

Topics:

- Modify Deployment Details and Configuration
- Start and Stop Service Manager and Deployments
- View and Edit the Configuration for Microservices
- Create and Enable the Authorization Profile for External Identity Provider

3.1.1 Modify Deployment Details and Configuration

Click the Deployment name in the Service Manager Overview page to open the Deployment Information page. Use the Deployment Information page to review, modify, and manage the selected service configuration.

Details Tab

Use to review the selected deployment configuration. All the deployment directories that you configured with the Configuration Assistant are displayed.

For Oracle database, the only directory that you can edit is the Oracle GoldenGate home (`OGG_HOME`). This allows you to use a different installation than the one you originally configured.

For SQL Server and Db2 z/OS, you need to follow the steps given in the Setting up Environment Variables for Db2 z/OS , Setting up for DB2, and Setting up for SQL Server in the *Using Oracle GoldenGate on Oracle Cloud Marketplace* guide.

> **Note:**
>
> It's important to do the settings for SQL Server and Db2 z/OS to make sure that the Administration Service starts when using either of these databases.

Configuration Tab

Use to add, edit, or delete the environment variables for your deployment.

For Oracle database, check if the `TNS_ADMIN` environment variable is added to this list. Click the plus sign next to the **Environment Variables** to add `TNS_ADMIN`. This variable defines the location of the `$OGG_Deployment/etc` directory, where the `tnsnames.ora` file is added and stored. You need the `tnsnames.ora` file to set up your database connections while adding database credentials for Oracle database.

Here's a sample structure of the `tnsnames.ora` file for a multitenant container database.

```
# tnsnames.ora Network Configuration

LISTENER_ORCL =
  (ADDRESS = (PROTOCOL = TCP)(HOST = main.oracle.com)(PORT = 1521))


CGGNORTH =
  (DESCRIPTION =
   (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP)(HOST = main.us.oracle.com)(PORT = 
1521))
   )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = orcl.us.oracle.com)
    )
)
```

Certificates

Use this tab to manage certificates for the server, client and CA certificates. See Add and Manage Certificates for a Deployment for details.

Authorization Profiles

Use this tab to delegate user and group management to third party ID providers such as Oracle Identity Cloud Service. Integration with an external Identity Management (IDM) system using OpenID/OAuth2.0 protocol provides Oracle GoldenGate users with:

- A single sign-on experience
- Ease of deploying Oracle GoldenGate cloud integration with IDCS.

See Create and Enable the Authorization Profile for External Identity Provider for details.

3.1.2 Start and Stop Service Manager and Deployments

Starting and Stopping the Service Manager

The start and stop process of the Service Manager within Oracle GoldenGate Microservices Architecture is different based on how the Service Manager is configured within your environment.

- If the Service Manager is configured in manual mode then there are scripts in the `$DEPLOYMENT_HOME/servicemanager/bin` directory that you can run to start or stop the Service Manager. The `$DEPLOYMENT_HOME` is the directory where Oracle GoldenGate is installed.
  - To start the Service Manager: `$DEPLOYMENT_HOME/servicemanager/bin/startSM.sh`

  – To stop the Service Manager: `$DEPLOYMENT_HOME/servicemanager/bin/stopSM.sh`

> **Note:**
>
> If you want to start or stop the Service Manager, you also have to set the `$OGG_ETC_HOME` and `$OGG_VAR_HOME` to the Service Manager sub-directories.

- If the Service Manager is configured as a daemon, the scripts required to start or stop for manual interaction are not created. The operating system is responsible for starting or stopping the Service Manager.

  For OEL 7 and OEL 8:

  ```
  systemctl start OracleGoldenGate
  ```

  ```
  systemctl status OracleGoldenGate
  ```

  ```
  systemctl stop OracleGoldenGate
  ```

- If the Service Manager is configured to run with the XAG agent in an Oracle Cluster Ready Service (CRS); then the start and stop process is handled by the CRS stack.

Stopping and Starting Deployments and Other Microservices

> **Note:**
>
> If Oracle GoldenGate Service Manager is registered as a system daemon, then the Service Manager along with the other servers, are automatically started when the host is (re)started.

1. Log in to your Service Manager instance as the system adminstrator.
2. In the **Deployments** section of the Service Manager home page, locate the deployment that you need to start or stop.
3. In the **Actions** column, click **Start**.
4. Verify if all the services associated with the deployment have started. An indication that the services have started is that the **Action** column automatically shows the **Stop** option. By default, all server instances are in **Running** state after the deployment process is complete.
5. To start or stop a service, such as the Administration Service or the Distribution Service, go to the Services section.
6. Identify the server (or service) that you need to start (or stop) and click start in the **Action** column, the same way you did for Deployments.

3.1.3 View and Edit the Configuration for Microservices

Use the Service Manager Overview page to view and edit the configuration and restart options for Administration Service, Distribution Service, Performance Metrics Service, and Receiver Service.

Click the **Details** settings icon for the microservice to check the service configuration. The Service Configuration page is displayed. This page allows you to view and edit the configuration and the restart options for the corresponding service. The configuration and restart options for all the services are the same.

The following table explains the Service Configuration and Restart Options on the Services Configuration page.

Service Configuration Options	**Description**
Port	Port Number for the corresponding service
U-Mask	File mode creation mask
Enabled	Indicates that the service is managed by Service Manager.
Status	Indicates the status of the service.
Restart Options	**Description**
Enabled	If set to true, then the service will attempt to restart automatically if it encounters an error.
On Success	If set to false, then the service is only restarted if it fails.
Delay	The time (in minutes) to pause between discovering that a process is terminated abruptly and restarting it.
Retries	The maximum number of trials to restart the service, before aborting the retry effort.
Window	The time interval in which the retries are counted. The default is 120 minutes.
Disable on Failure	If set to true, the service is disabled after it fails all execution attempts in an execution window.

3.1.4 Create and Enable the Authorization Profile for External Identity Provider

To access the Authorization Profiles page, click the deployment name or the Service Manager name from the **Service Manager Overview page's Deployment** section.

From the Deployment or Service Manager Details page, click the **Authorization Profiles** tab to delegate authentication and authorization to third party ID Providers such as Oracle Identity Cloud Service (IDCS).

Select the plus sign (+) next to the Profiles section start creating an authorization profile. Enter the following details for the profile:

- Profile Name: Name of the authorization profile.
- Description (optional): Short summary of the profile being created.
- Enable Profile: Activates the profile for the deployment.
- Authorization Profile Type: IDCS
- Tenant Discovery URI: IDP server's OpenID Discovery Docs endpoint (`/.well-known/openid-configuration`).

- Client ID: IDP application’s client ID.
- Client Secret: IDP application’s client secret (securely stored).

In the Group Mapping section, the user mapping for IDCS groups to Oracle GoldenGate user roles is configured. You need to enter the name of the IDCS group with the corresponding user role. These values are case-sensitive.

Here are the user role options that map the name of a group with respective role in IDCS:

- Security Role
- Administrator Role
- Operator Role:
- User Role

Click **Submit** to create an authorization profile.

You can enable the authorization profile as soon as you create it or you can enable it later from this **Deployment** details page. To enable the authorization profile, select the authorization profile that you want to enable and click the **Enable Profile** toggle switch.

3.2 Enable and Use Debug Logging

You can enable debug logging and download debug log files from this page.

Enabling Debug Logging:

To enable debug logging:

1. Click the **Debug Log** option from the navigation pane of the Service Manager page.
2. Click the **Enable Debug Log** toggle switch to start logging debug information.

Using the Debug Log

You can access and use the debug log file from this page:

1. Click the **Download Debug Log File** option to save a local copy of the debug log
2. Click the **Load Debug Log File** option to view the debug log on this page.
3. Click the **Delete Debug Log File** button to delete a debug log.
4. Search for specific entries in the debug log using the **Search By** box, if required.
5. You can click **Refresh** to get the latest log information, if it doesn't get refreshed automatically.

.

3.3 Read the Log Information

You can review all of the messages logged for your Service Manager with this page.

Using the Table

An updated log of connected distribution path and target initiated paths is displayed. You can sort the list by date or severity by clicking on the adjacent arrow. Also, you can refresh this log and choose how many pages you want to view.

To search, you select Date, Severity, or Message, and then select the appropriate options to construct your search.

Notice the **Notifications** tab at the bottom of the page. It displays messages from the service, which are not updated in the log due to transaction errors. For example, failure to log in to the database using the database credentials.

3.4 Add and Manage Certificates for a Deployment

You can either manage the certificates from a specific deployment or Service Manager. These certificates are shared because there might be multiple deployments supervised by one Service Manager. Within an individual deployment, you have the choice of deciding to share a certficate between deployments or to keep it local.

Click the **Certificate Management** tab from the left navigation pane of the Service Manager. Select the deployment from the drop down list to view information about the server, client certificates and CA certificates. The time period of validity with the used signing algorithms from the issuer are displayed.

Click the **Detail** icon from the **Action** column of the certificate store table to view details about the certificate including issuer of the certificate, target for the certificate, and signature algorithm.

Click the **Replace** (pencil) icon to replace server certificates.

> **Note:**
>
> You cannot modify or edit an existing certificate. You can only replace it with a new certificate.

Click the **Delete** icon in the Action column to delete the certificate.

Add Client Certificate

To add a client certificate:

1. Click the plus (+) sign next to the Client Certificates section. Add Client Certificate dialog box appears.
2. Enter the following details for the client certificate:
   - Unique Name: Name of the certificate.
   - Certificate PEM: Enter a certificate .pem file or upload a .pem file.
   - Private-Key PEM: Enter or upload the private key for the .pem file.
   - CA Certificates: Enter or upload the CA certificate.
3. Click **Add**.

Add CA Certificate

To add a CA certificate:

1. Click the plus (+) sign next to CA Certificates. Add CA Certificate dialog box appears.
2. Enter the following details for the CA certificate:

- Unique Name for the CA certificate.
- Certificate PEM value can be entered in the box or uploaded.
- Certificate location can be shared. CA Certificates for the Service Manager are always shared and cannot be local. When adding or replacing CA certificates, the **Shared** option is always force-checked.

3. Click **Add**.

Also see Secure Deployment Requirements and Two Deployments: Certificate Requirements in *Step by Step Data Replication Using Oracle GoldenGate Microservices*.

3.5 Add Users from the Service Manager

Each deployment has its own list of users, and when you add users, you add them to that deployment.

You can create users from the Service Manager or the Administration Service. See Add Oracle GoldenGate Deployment Users from the Administration Service for steps to create users.

The only user that can manage the services in Service Manager is the user that was originally added as the security user when you initially added the deployment to the Service Manager using `oggca`. The other users are specific to the MA deployment and the security user needs to create users to every MA deployment individually.

You can create users for that deployment by performing the following steps:

1. Log in to either the Service Manager or the Administration Service.
2. From the left navigation pane, select Administrator.
3. Click Users (+) to add users.
4. Enter a unique user name.
5. Select one of these roles from the Role list box:

Role ID	Privilege Level
User	Allows information-only service requests, which do not alter or effect the operation of either the MA. Examples of Query/Read-Only information include performance metric information and resource status and monitoring information.
Operator	Allows users to perform only operational actions, such as creating, starting and stopping resources. Operators cannot alter the operational parameters or profiles of the MA server.
Administrator	Grants full access to the user, including the ability to alter general, non-security related operational parameters and profiles of the server.
Security	Grants administration of security related objects and invoke security related service requests. This role has full privileges.

6. Select the user type from the Type list box as **Basic** (digest authentication) or **Certificate**.
7. Enter information that describes the user.
8. If you select the user type as Basic, then the authentication is done based on the username and password. So, the **Password** option comes up, if you select the **Basic** security type and not with the **Certificate** option. Enter the password twice to verify it.

   If you select the user type as Certificate, then the user will autherntticate themselves by presenting a client certificate. After you select the Certificate option, you need to enter the common name (in the certificate that will be presented such **CN="certuser"**.

   > **Note:**
   >
   > The certificate is with the user and not saved by the Oracle GoldenGate service. When presented for autherntication, the Oracle GoldenGate service first authenticates that the certificate presented can be trusted and then checks to see that the common name in the certificate has been registered as a valid user. If yes, it will assign the appropriate user role.

9. Click **Submit**.

   The user is registered

Users cannot be changed. You must delete a user, and then add it again. However, you can modify or edit a user's attributes, by clicking the **Edit User** (pencil) in the Action column of the **Users** table.

You can switch the User Type from Basic to Certificate or the other way around.

You can also change the password for the user, if required.

Click **Submit** to confirm the modifications to the user attributes.

4

Configure Data Replication Processes from the Administration Service

You can perform all Extract, Replicat, and database credential setup tasks from the Administration Service.

The Administration Service contains options to:

- Add table and schema level transaction logging (`TRANDATA`/`SCHEMATRANDATA`)
- Add checkpoint and heartbeat tables
- Create role-based users that are authorized to do more granular tasks than that of the security user that gets created when adding a deployment
- Configure autostart and restart processes
- Create and manage encryption profiles

Unlike users that are created in the Service Manager deployment, Administration Service users can create Extracts, Replicats, Paths, Credentials, and adjust deployment related settings, while Service Manager users can enable, disable, start and stop deployments and individual services.

You can also create users with the **Admin** role from the Administration Service although you can create users from the Service Manager as well.

The deployment users are created from the Service Manager web interface. These users can start and stop microservices and the Service Manager itself.

Users created from the Administration Service can create Extract, Replicat, and other processes.

Topics:

4.1 Quick Tour of the Administration Service Overview Page

You can access the Administration Service from the Service Manager Overview page or by directly specifying the URL in the web browser.

Use the Administration Service to add, and immediately start managing and monitoring Extract and Replicat processes.

The table on the home page displays the severity of critical events. Use the left-navigation pane to access various configuration details, a list of severity issues with their diagnosis.

Here are some of the key actions that you can perform from this page.

Action	Description
View the home page in tabular format	Use the Table Layout swivel to turn the tabular format on and off.
View Extracts and Replicats	The statistical representation the home page displays current state of Extracts and Replicats (Starting, Running, Stopped, Abended, Killed)
Add an Extract	See How to Add an Extract
Add a Replicat	See How to Add a Replicat
Stop and start Extracts	Using Extract Actions
Stop and start Replicats	See Using Replicat Actions
View and search critical events	Monitor severity of events using the Critical Events table and also search for specific events, if required.

- Review Critical Events
  You can review and search for critical events from the Review Critical Events section of the Administration Service home page.

4.1.1 Review Critical Events

You can review and search for critical events from the Review Critical Events section of the Administration Service home page.

Once you set up the Extracts and Replicats along with the distribution path, you are able to see the critical events associated with them.

Search for Critical Events from the Review Critical Events Table

The Review Critical Events table displays the severity, error code, and error messages for critical events. You can view 20 error messages on a single page and you can also search for specific events.

Additionally, you can examine events in depth from the Performance Metrics Service. For details see Quick Tour of the Performance Metric Server home page.

4.2 Configurations from the Administration Service

The Configuration page can be accessed by clicking the Configuration option from the left-navigation menu of the Administration Service. Use this Configuration page to perform the following tasks:

Topics:

- Add and Alter Database Credentials

- Set up and Use the Master Keys and Encryption Keys
- Access the Parameter Files
- Automate Maintenance Tasks

4.2.1 Add and Alter Database Credentials

To create and run Extract and Replicat processes, you need to set up database credentials.

1. Launch the Administration Service interface and log in.
2. Click **Configuration** from the **Application Navigation** pane.
3. Click the plus sign (**+**) sign next to Credentials.
4. Enter the following details in the displayed fields:

Database Credential Options	Description
Credential Domain	Specify a domain name to which the database credential is associated. For example, "OracleGoldenGate" is the default domain name, incase you don't specify a domain name.
Credential Alias	This is the alias for your database credential.
User ID	This is the username of the database user. For Oracle database, if you use the EZconnect syntax to connect to the database, then you can specify the value in this field in the following manner: *dbusername*@*hostname*:`port/service_name` `dbusername` is the database user name. `hostname` or IP address of the server where the database is running. `port` is the port number for connecting to the database server. Usually, this value is 1521. `service_name` is the name of the service provided in the tnsnames.ora file for the database connection.
Password	Password used by database user to log in to the database.

5. Click **Submit**.
6. Click the Connect to database icon to test that the connection is working correctly. If the connection is successful, the Connect to database icon turns blue.

When you successfully log into your database, you can add and manage checkpoint tables, transaction information (TRANDATA), and heartbeat tables. All of the tables can be searched using the various search fields. As you type, the table is filtered and you can use the search button with the search text.

Alter Database Credentials

1. From the Action column of the database credential that needs to be altered, select the pen icon.
2. You can edit the the user ID and password for the credential.

3. Click the trash icon from the Action column, to delete a database credential.

- Configure Kerberos Authentication with MA

4.2.1.1 Configure Kerberos Authentication with MA

Here are the steps to configure kerberos authentication from the Admin Client.

Connect to the Administration Service from the Admin Client:

```
CONNECT http://localhost:9005 DEPLOYMENT oggdep as ggadmin PASSWORD
We1come_$
```

Alter the credentialstore after connecting to the Administration Service of the deployment `oggdep`:

```
ALTER CREDENTIALSTORE ADD USER /@DBEAST NOPASSWORD ALIAS ggeast
```

Output shows:

```
2020-06-22T21:08:33Z  INFO OGG-15102  Credential store created.
2020-06-22T21:08:33Z  INFO OGG-15114  Credential store altered.
```

Run the following command to verify that the credentialstore was altered successfully:

```
INFO CREDENTIALSTORE
```

Output displays the following:

```
Default domain: OracleGoldenGate
  Alias: ggeast
  Userid: /@DBEAST
```

When using the MA web UI to create the credential, if the **User ID** field begins with a ***/*** character, then the password is not required. So, in the **User ID** field, enter / *connect_string* where *connect_string* is your connection string.

Here, the `NET SERVICE` is the simple name for the database service. Alternatively, a complete connect string (descriptor) can be used instead of the Oracle net service name.

Here's an example of a predefined net service name and connect descriptor mapping:

```
DBEAST = (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=db1))
(CONNECT_DATA=(SERVICE_NAME=DBEAST.regress.rdbms.test.example.com)))
```

- Example: Using USERIDALIAS in Parameter File for Kerberos Account

4.2.1.1.1 Example: Using USERIDALIAS in Parameter File for Kerberos Account

The following example shows how to set the `USERIDALIAS` values in the parameter file after creating the credential store with Kerberos authentication:

```
ALTER CREDENTIALSTORE ADD USER /@ggadmin NOPASSWORD ALIAS ggadmin



2020-12-17T21:08:33
INFO    OGG-15102  Credential store created.2020-12-17T21:08:33
INFO    OGG-15114  Credential store altered.


ALTER CREDENTIALSTORE ADD USER /@ggadmin_mining NOPASSWORD ALIAS
ggadmin_mining


2020-12-17T21:09:45
INFO    OGG-15102  Credential store created.2020-12-17T21:09:45
INFO    OGG-15114  Credential store altered.


INFO CREDENTIALSTORE



Default domain: OracleGoldenGate
Alias: ggadmin
Userid: /@ggadmin

Default domain: OracleGoldenGate
Alias: ggadmin_mining
Userid: /@ggadmin_mining
```

After altering the credentialstore, you can specify `USERIDALIAS` options in the parameter file:

```
USERIDALIAS ggadmin
DOMAIN OracleGoldenGate
TRANLOGOPTIONS MININUSERIDLIAS ggadmin_mining
DOMAIN OracleGoldenGate
```

4.2.2 Set up and Use the Master Keys and Encryption Keys

You can set the master keys and encryption keys using the **Key Management** tab in the **Configuration** page of the Administration Service.

Using Master Keys

If you want to encrypt your data, then create a Master Key by clicking the + sign in the Master Key section. The master key is generated automatically.

You can change the status of the key to Available or Unavailable, by clicking the edit icon in the Master Key table. You can also delete the Master Key from the table by clicking the delete icon.

For details on the Master Key concept, see Encrypting Data with the Master Key and Wallet Method. .

Using the Encryption Keys

To use this method of data encryption, you configure Oracle GoldenGate to generate an encryption key and store the key in a local `ENCKEYS` file. The `ENCKEYS` file must be secured through the normal method of assigning file permissions in the operating system. This procedure generates an AES encryption key and provides instructions for storing it in the `ENCKEYS` file.

To generate the ENCKEYS files, click the + sign in the Encryption Keys section. The Encryption Keys is generated.

For details on the Encryption Keys concept, see the Encrypting the Data with the ENCKEYS Method.

4.2.3 Access the Parameter Files

View and configure the Extract, Replicat, and Global parameter files from the **Parameter Files** tab of the **Configuration** page of the Administration Service.

To use the different parameter file options:

1. Select the **Configuration** option from the Administration Service left-navigation pane.
2. Select the **Parameter Files** tab.

   A list of existing parameter files is displayed along with the GLOBALS parameter file.
3. If you select any of the parameter files, you are presented with the option to edit or delete the selected file. If you want to change the GLOBALS parameter file, you need to restart the Administration Service and any Extracts and Replicats..
4. Click + add parameter files.
5. Enter the file name and the required parameters. Make sure to enter the file name with the `.prm` extension.
6. Click **Submit**. The new parameter file is displayed in the list of parameter files.

The actual location of the parameter files on the disk can be determined using the following step:

1. Identify the GoldenGate Deployment ETC Home:
   a. Go to Service Manager Overview page.
   b. Click the deployment from the Deployments section for which you need to find the parameters file.
   c. Under the Deployment Detail window, navigate to the Oracle GoldenGate deployment `/etc` home directory.
   d. Go into the `/config/ogg` directory where the parameter file is located.

The following example shows how to navigate to your parameter file location:

```
[oracle ~]$ cd /opt/app/oracle/gg_deployments/oggdepl/etc
[oracle etc]$ cd conf/ogg[oracle ogg]$ lsEXT_DEMO.prm GLOBALS REP_DEMO.prm
```

4.2.4 Automate Maintenance Tasks

Use the **Tasks** tab on the Configuration page, to set up the following automated tasks.

Purging Trails

The Purge Trail page works the same way as the Manager `PURGEOLDEXTRACTS` parameter in the Classic Architecture. It allows you to purge trail files when Oracle GoldenGate has finished processing them. Automating this task ensures that the trail files are periodically deleted to avoid excessive consumption of disk space.

From the Tasks tab, when you select the Purge Trail page, it allows you to configure the Administration Service purge trail process.

1. Add a Purge Trail task by clicking the + sign .
2. Enter the **Operation Name** of the Administration Service task. The operation name is case sensitive. For example, you can create an operation with the name **TASK1** and another operation named **task1**.
3. Enter the trail path or trail name in the **Trail** field.
4. Click the **+** sign to add the trail to the **Selected Trails** list.
5. If you don’t need to use checkpoints, disable the option **Use Checkpoints**. However, Oracle recommends using checkpoints. If you don't use checkpoints. the trail will be purged whether or not it has been consumed if the keep rule is met.
6. Set the **Keep Rule** value to specify the maximum number of hours, days, or number of files for which the Purge Trails task needs to be active.
7. Specify the number of hours or days when the purge trails task has to run, in the Purge Frequency field and click Submit.
8. Use the Purge Trails task table to edit or delete the task, as required.

   Also see PURGE EXTTRAIL.

Purging Tasks

You can automatically purge processes associated with an Administration Service.

From the Tasks tab, click Purge Tasks.

1. Enter the **Operation Name** that you need to set up for automatic purging.
2. Select the Extract or Replicat task (initial load process) **Process Name** for the operation. The list contains all processes so ensure that you select the correct task.
3. Select the Extract or Replicat task (initial load) **Process Type** for the operation.
4. If you enable **Use Stop Status**, the status of the task is used to perform the purge task.
5. Enter the hours or days after which you need to purge the process and click **Submit**.
6. Edit or delete the purge process task using the relevant icon from the Purge Tasks table.

Reporting Lag

You can manage lag reports from the Lag Report tab. To do so:

1. From the Tasks tab, click Lag Report.
2. The Action column contains all the options to delete, alter, refresh, and view the lag report task details.
3. Select the required option.
4. If you select the Alter Task option, you are presented with options to edit the lag report. The options are:
   - Enabled: To keep processing the lag report task.
   - Check Every (in minutes): To set a time interval to check the lag report.
   - Report: To log report for the task.
   - If Exceeds: To specify a threshold after which a warning would be initiated.
   - Warning: To allow a warning to be generated incase the lag threshold exceeds the specified limit.
   - When Exceeds: The lag threshold after which the warning is triggered.
5. Click Submit.

4.3 Add and Manage Extract

The Extract process is used to set up the data capture mechanism for Oracle GoldenGate.

Topics:

- Before Adding an Extract
- Add an Extract

4.3.1 Before Adding an Extract

Before performing the tasks in this topic, make sure that you are able to connect to the database from the web interface. See Add and Alter Database Credentials for steps to add database credentials and test the connection to the database.

After you have connected to the database, the Configuration page will show the checkpoint and heartbeat configuration sections.

Enable TRANDATA or SCHEMATRANDATA Information

Valid for Oracle and heterogeneous databases.

Depending on the source database, supplemental logging must be enabled. This can be done at the table, schema, or global (database) level.

To enable supplemental logging at the table and schema level, on Configuration page:

1. Select the **Table** or **Schema** option as required and click plus sign to add.

2. Enter the name of the table for which you need to set up supplemental logging. Make sure to enter the full table name with schema name, such as, `schema.table1`. You can also use wildcard instead of specific table name.
3. Select the **Add TRANDATA Information in the background?** option as required.
4. Click **Submit**.

You can also use the commands `ADD TRANDATA` and `ADD SCHEMATRANDATA` for setting up trandata and schema level trandata. For details, see `ADD TRANDATA` and `ADD SCHEMATRANDATA`. You can skip `ADD TRANDATA` in case of initial load without CDC.

Create Heartbeat Table

To create the heartbeat table, you have to follow these steps on the source and target system:

> **Note:**
>
> Creating the heartbeat table is optional but is recommended.

1. From the Administration Service, select **Configuration** from the navigation pane.
2. Select the + sign next to the Heartbeat section of the Database tab. You'll need to enter the values for the heartbeat frequency, retention time, and purge frequency.

You can create the heartbeat table using the `ADD HEARTBEATTABLE` command from the Admin Client or GGSCI. See `ADD HEARTBEATTABLE`.

Create the Oracle GoldenGate CDC Cleanup Task

For SQL Server users, there is a requirement to create Oracle GoldenGate CDC Cleanup tasks before adding an Extract. You can do so by performing the steps in Details of the Oracle GoldenGate CDC Cleanup Process

in the *Using Oracle GoldenGate for Heterogeneous Databases* guide.

4.3.2 Add an Extract

Set up database credentials to create and run Extract using the steps in Add and Alter Database Credentials.

Now, you're ready to add an Extract for your deployment.

1. From the Overview page of the Administration Service, click the + sign next to Extracts.
2. Choose the type of Extract to create and click **Next**.

   > **Note:**
   >
   > To learn about creating Initial Load Extract, see Loading Data from File to Replicat in MA in *Administering Oracle GoldenGate*.

   You can also create a Change Data Capture Extract for MySQL and SQL Server databases.

3. Provide the required information designated with an asterisk (*). Here's a description of the options in the different sections for the Add Extract screen:

Option	Description	Database
Basic Information Section		
Process Name	Name of the Extract process. The name of the Extract process can be up to 8 characters.	All databases
Description	Description of the Extract process being created.	All databases
Intent	Describes the purpose of creating the Extract. The default option is Unidirectional. Other options are High Availability, Disaster Recovery, N-Way, which are informational only.	All databases
Begin	Used to set the beginning location in the redo or transaction log from which the Extract will start to capture data. Available options are Now, Custom Time, CSN or Position in Log, and EOF depending on the supported database.	All databases
Trail Name	A two character trail name.	All databases
Trail Subdirectory, Size, Sequence, and Offset	You can further configure the trail details.	All databases
Remote	Enable this option if the Extract trail is remote. For Oracle databases, enable this option if the Extract trail is to be written directly to a remote Oracle GoldenGate Classic installation. For MySQL, setting this option enables the `TRANLOGOPTIONS ALTLOGDEST REMOTE` parameter to support a remote Extract, and is not related to trails.	Oracle, MySQL
Registration Information Section		
CSN	Commit Sequence Number (CSN) value	Oracle

Option	Description	Database
Share	Choose the method to share the LogMiner data dictionary. Options are: • Automatic: This option allows the system to choose the method for sharing the dictionary . • None: Choosing this option, will not allow the dictionary to be shared. • Extract: Choose this option to allow sharing the logminer dictionary for specific Extract.	Oracle
Optimized	Enable this option to optimize the Extract registration.	Oracle
Downstream Capture	Enable this option to set up a downstream Extract for log mining.	Oracle
Register Only	Use this option to just register the Extract and not add the Extract. The registration creates the replication slot when you register the Extract or use the Register Only option.	PostgreSQL
Source Database Credential		
Create new credential	If you haven't set up your database login credentials, you can create and save the database login credentials from here.	All
Credential Domain	Create a domain for the database.	All
Credential Alias	Specifiy a credential for the database login.	All
User ID	Specify a user name for logging into the database.	All
Password, Verify Password	Enter the password used to login to the database and reenter the password to verify.	All
Credential Domain	Saves the credential user under the specified domain name. Enables the same alias to be used by multiple Oracle GoldenGate installations that use the same credential store. The default domain is Oracle GoldenGate.	All databases

Option	Description	Database
Credential Alias	Specifies an alias for the user name. Use this option if you do not want the user name to be in a parameter file or command. If `ALIAS` is not used, the alias defaults to the user name, which then must be used in parameter files and commands where a login is required. You can create multiple entries for a user, each with a different alias, by using the `ADD USER` option with `ALIAS`.	All databases
Downstream Mining		
Mining Credential Domain	Domain name of the downstream mining database.	Oracle
Mining Credential Alias	Alias for the mining downstream database.	Oracle
No UserID	Enable this option if there is no source database connection. Selecting this option enables the ADG fetch options.	Oracle
ADG Fetch Credential Domain	Domain name for the ADG fetch database.	Oracle
ADG Fetch Credential Alias	Domain alias for the ADG fetch database.	Oracle

4. (Optional) Enter the encryption profile description. If you have not created an encryption profile, then the Local Wallet profile would be selected by default. :
   a. Select the profile name from the list box. You can select the Local Wallet or a custom profile.
   b. Select the encryption profile type from the list box.
   c. Specify the masterkey for the encryption profile. This option doesn't exist with SQL Server.
5. This is an optional step. Enter the Managed Options while creating all types of Extract processes. The following table provides these options:

Option	Description
Profile Name	Provides the name of the autostart and autorestart profile. You can select the default or custom options. If you have already created a profile, then you can select that profile also. If you select the Custom option, then you can set up a new profile from this section itself.

Option	Description
Critical to deployment health	(Oracle only) Enable this option if the profile is critical for the deployment health. **Note:** This option only appears while creating the Extract or Replicat and not when you set up the managed processes in the Profiles page.
Auto Start	Enables autostart for the process.
Startup Delay	Time to wait in seconds before starting the process
Auto Restart	Configures how to restart the process if it terminates
Max Retries	Specify the maximum number of retries to try to start the process
Retry Delay	Delay time in trying to start the process
Retries Window	The duration interval to try to start the process
Restart on Failure only	If true the task is only restarted if it failes
Disable Task After Retries Exhausted	If true then the task is disabled after exhausting all attempts to restart the process.

6. Click **Next**.
7. You can edit the parameter file in the text area to list the table details that you are interested in capturing. Here's a sample of the Extract parameter file:

   ```
   EXTRACT exte
   USERIDALIAS cggnorth DOMAIN OracleGoldenGate
   EXTTRAIL east/ea
   UPDATERECORDFORMAT COMPACT
   SOURCECATALOG DBEAST
   DDL INCLUDE MAPPED OBJNAME hr.*
   DDLOPTIONS REPORT
   TABLE hr.*;
   ```

   To know more about parameter files, see Using Oracle GoldenGate Parameter Files

   .
8. You can select **Register Extract in the background** to register the Extract in the background asynchronously. This option is required for Oracle and PostgreSQL databases.
9. Click **Create and Run** to create and start the Extract. If you select **Create**, the Extract is created but you need to start it using the Extract drop-down on the Overview page.

   You are returned to the Overview page of the Administration Service.

Select the **Action** list if you want to look at the Extract details such as process information, checkpoint, statistics, parameters, and report.

Topics:

- Add an Initial Load Extract
- Using Extract Actions

4.3.2.1 Add an Initial Load Extract

An initial load Extract pulls data from tables and writes the records to an external file (`EXTFILE`) rather than to a trail (`EXTTRAIL`).

Common uses for an initial load Extract are to instantiate the data to a non-Oracle target, such as from Oracle to SQL Server or from MySQL to DB2.

See Instantiating Oracle GoldenGate with an Initial Load in Microservices Architecture to use cURL for performing this procedure.

Before You Begin

Verify the following before you start adding an initial load Extract:

- Check that the environment variables are set correctly for your deployment.
- Disable the target constraints and truncate tables.

  Here an example of disabling constraints of tables:

  ```
  ALTER TABLE hr.countries DISABLE CONSTRAINT COUNTR_REG_FK;
  ALTER TABLE hr.departments DISABLE CONSTRAINT DEPT_LOC_FK;
  ```

  Do this for all tables that need to be loaded on the target database using initial load.

  Here's an example of truncating target tables:

  ```
  TRUNCATE TABLE hr.job_history;
  TRUNCATE TABLE hr.employees;
  ```

- Verify the database connection for source and target databases.

Create the Change Data Extract

Here are the steps to add a Change Data Extract:

1. Log into the Administration Service.
2. Click the plus sign in the **Extracts** section.
3. Select Integrated Extract and click **Next**.
4. Enter the Extract options and click **Next**. Make sure to select the Register to PDB option if you are connecting to Oracle database.

5. Check the Extract parameter options. The Extract parameter file looks similar to this:

   ```
   EXTRACT exte
   USERIDALIAS ggeast DOMAIN OracleGoldenGate
   EXTFILE ea
   TABLE EAST.HR.*;
   ```

6. Click **Create and Run** to start the Extract.

Create the Change Replicat

Ensure that you don't start this Replicat after creating it. Here are the steps to set up the Replicat:

1. Click the plus sign in the Replicats section of the Administration Service Overview page.
2. Select Nonintegrated Parallel Replicat and click Next.
3. Specify the Replicat Options and select the precreated checkpoint table or create one if not done previously.
4. Review the Replicat parameter file, which would look similar to this:

   ```
   REPLICAT repe
   USERIDALIAS ggwest DOMAIN OracleGoldenGate
   MAP EAST.HR.*, TARGET HR.*;
   ```

5. Click **Create** to add the Replicat. *Do not click* **Create and Run**.

Determine the SCN Value of the Transaction

Determine the SCN of the intial position of the open transaction to set up the initial load Extract. If there are no existing transactions, then the current SCN is used. Here are the steps to determine the SCN from the source database:

1. Log into the source database.
2. Run a query similar to the following, to view active transactions:

   ```
   Select T.START_SCN, T.STATUS TSTATUS, T.START_DATE,
          S.SID, S.SERIAL#, S.INST_ID, S.USERNAME, S.OSUSER, S.STATUS 
   SSTATUS, S.LOGON_TIME
          From gv$transaction T
          Inner
            Join gv$session S on S.SADDR = T.SES_ADDR
          Union All
   ```

   The query to determine the current status of the transaction is similar to the following:

   ```
   Select current_scn, 'CURRENT', CURRENT_DATE,
          NULL, NULL, NULL, 'SYS', NULL, NULL, NULL from v$database
          Order by 1;
   ```

3. From the TSTATUS column displayed in the output, you can determine the ACTIVE SCN or the CURRENT SCN value, if the ACTIVE SCN value is not available.

Create the Initial Load Extract

If there is an active open transaction, then use the start SCN from that transaction. If not, you can use the current SCN value. The following instructions assume that there is an active transaction, so the start SCN value is used. Here are the steps for setting up the initial load Extract:

1. Click the plus sign in the Extracts section of the Administration Service Overview page.
2. Select Intial Load Extract and click **Next**.
3. Provide the Extract name, for example, `exteinit`, and click Next.
4. Check and modify the parameters for this Extract. The Extract parameters need to include the SQLPREDICATE parameter, as shown in the following example:

   ```
   EXTRACT EXTEINIT
   USERIDALIAS ggeast DOMAIN OracleGoldenGate
   EXTFILE ei MEGABYTES 250 PURGE
   TABLEEXCLUDE EAST.HR.EMP_DETAILS_VIEW
   TABLE EAST.HR.*; SQLPREDICATE "AS OF SCN 2898620";
   ```

5. Click **Create and Run** to start the intial load Extract. This type of Extract stops after it finishes extracting all data from the tables in the HR schema and writes to the trail file.
6. To check that the data extraction happened accurately, access the report file using the **Action** button of the intial load Extract. Select **Details** and then **Report**. It

Create the Initial Load Replicat

Here are the steps to create an initial load Replicat:

1. Click the plus sign from the Replicats section of the Administration Service Overview page.
2. Select **Integrated Replicat** and click **Next**.
3. Specify the Replicat options including the database credentials domain of the target database, and the checkpoint table name, and click **Next**.
4. Check and modify the Replicat parameter file. Here's an example of the Replicat parameters:

   ```
   REPLICAT REPEINIT
   USERIDALIAS ggwest DOMAIN OracleGoldenGate
   MAP EAST.HR.*, TARGET HR.*;
   ```

5. Click **Create and Run**.
6. From the Administration Service Overview Page, click the Action button for the initial load Replicat, and click **Details**.
7. Click the **Statistics** tab to verify that the data has been loaded on the target database. It displays the details of tables that are loaded from source to target. You can check the Inserts, Updates, and Deletes, to be sure of the loaded transactions.

8. Enable the constraints in the target table that was disabled previously, before starting the change Replicat. Here's an example showing the command to enable the same constraints:

```
-- Enable Contraints after the Initial Load is completed
ALTER TABLE hr.countries ENABLE CONSTRAINT COUNTR_REG_FK;
ALTER TABLE hr.departments ENABLE CONSTRAINT DEPT_LOC_FK;
ALTER TABLE hr.employees ENABLE CONSTRAINT EMP_JOB_FK;
ALTER TABLE HR.job_history ENABLE CONSTRAINT JHIST_JOB_FK;
ALTER TABLE hr.departments ENABLE CONSTRAINT DEPT_MGR_FK;
ALTER TABLE hr.employees ENABLE CONSTRAINT EMP_MANAGER_FK;
ALTER TABLE hr.job_history ENABLE CONSTRAINT JHIST_EMP_FK;
ALTER TABLE hr.locations ENABLE CONSTRAINT LOC_C_ID_FK;
ALTER TABLE hr.employees ENABLE CONSTRAINT EMP_DEPT_FK;
ALTER TABLE hr.job_history ENABLE CONSTRAINT JHIST_DEPT_FK;
```

Start Change Replicat

Here are the required steps to configure and start the change Replicat.

1. Choose the option to start the Replicat. There are different ways to start the change Replicat:

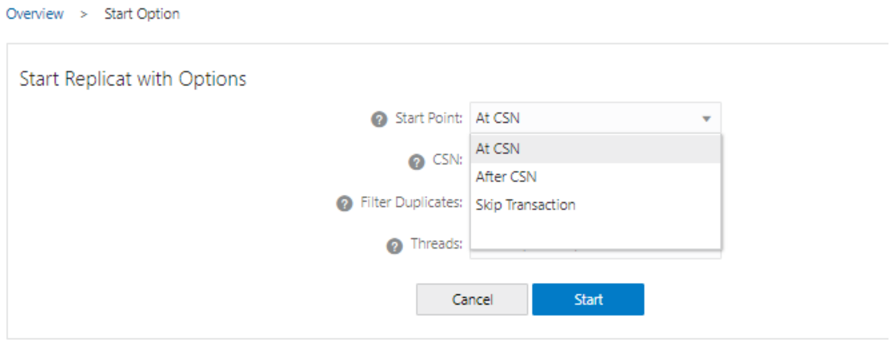

   - **At CSN**: If there are pending open transactions, which are not committed, then use the **At CSN** option and enter the start point, which is the active CSN value.
   - **After CSN**: If there are committed transactions after the recorded CSN value, then you can select the **After CSN** value.
   - **Skip Transaction**: Use the **Skip Transaction** option if you don't want to load data for the last open uncommitted transaction.

2. Navigate to the Administration Service Overview page and check the status of the Replicat. It should be running.

This completes the intial load set up for data replication.

Also see this blog Oracle GoldenGate Microservices Initial Load Instantiation with WebUI for steps to set up initial load for Oracle database from the Oracle GoldenGate Microservices web interface.

4.3.2.2 Using Extract Actions

Extract actions include tasks like monitoring details for the Extract, checkpoint details, DDL/DML statistics, cache manager statistics, and other details.

Use the **Action** button to start or stop the Extract or view and manage its details. When you select the **Action**, **Details** option for an Extract, you can perform the following tasks for it.

Action	Result
Details	Displays the following tabs: • **Process Information**: The status of the selected Extract process including the type, credentials, and trail details including trail name, trail subdirectory, trail sequence, and trail size. • **Checkpoint**: The checkpoint log name, path, timestamp, sequence, and offset value. You can monitor the input details, such as when starting, at recovery, and the current state. The checkpoint output values display the current checkpoint details. • **Statistics**: The active replication maps along with replication statistics based on the process type. You sort the lost to view the entire statistical data, daily, or hourly basis. • **Cache Manager Statistics**: Access the global statistics and object pool statistics information for the Extract process from this page. • **Parameters**: The parameters configured when the process was added. You can edit the parameters by clicking the pencil icon. Make sure that you apply your changes. • **Report**: A detailed report of the process including parameter settings and a log of the transactions. You could copy the report text and save it to a file so that you can share or archive it.
Start/Stop	The Extract starts or stops immediately.
Delete	Allows you to delete Extract after you stop it. This option only appears after you stop the Extract.

When you change the status, the list options change accordingly. As status are changing, the icons change to indicate the current and final status. The events are added to the Critical Events table. Additionally, progress pop-up notifications appear at the bottom of the page.

4.4 Add and Manage Replicat

The Replicat process receives the trail data and applies it to the database.

Topics:

- Before Adding a Replicat
  Before you start adding Replicat, create the checkpoint table.
- Add a Replicat

4.4.1 Before Adding a Replicat

Before you start adding Replicat, create the checkpoint table.

Add Checkpoint Table

After connecting to the database from the Credentials page of the Administration Service, you can create the checkpoint table. If you have an exisiting checkpoint table, you will be able to see it listed in the Checkpoint section.

1. Click the plus sign to enable adding a checkpoint table.
2. Add the checkpoint table name in the format `table.checkpoint_table_name`.
3. Click **Submit**. After the checkpoint is created, you'll be able to see in the list of checkpoint tables.

To perform this task from the command line, see `ADD CHECKPOINTTABLE` in the *Command Line Interface Reference for Oracle GoldenGate*.

4.4.2 Add a Replicat

After you've set up your database connections and verified it, you can add a Replicat for the deployment by following these steps:

1. Click the + sign next to Replicats on the Administration Service home page. The Add Replicat page is displayed.
2. Select a Replicat type and click **Next**.

   > **Note:**
   >
   > Some Replicat types are only available for certain databases. All Replicat types may not be applicable to your database.

   The types of Replicat are:

   - Integrated Replicat
   - Nonintegrated Replicat: This option is displayed with Oracle database.
   - Classic Replicat: This option is displayed with non-Oracle databases.
   - Coordinated Replicat
   - Parallel Replicat: If you select this option, then select an integrated or nonintegrated parallel Replicat. In the Replicat parameters for a parallel Replicat, you may need to include some basic parameters that work with parallel Replicat.

     To know more about different types of Replicat, see Choosing from Different Replicat Modes.
3. Enter the following details in the Replicat Options page.

Option	Description	Replicat Type
Process Name	The name of the Replicat process. For coordinated and parallel Replicats, the limit is 5 characters.	All
Description	Describes the Replicat process, something that can be entered that is longer than 5 or 8 characters to identify this process.	All
Intent	What you want the Replicat to be used for, such as High Availability or the Unidirectional default. The N-way option should be selected to save active-active configuration.	All
Source	Select the source to use, Trail or File.	All
Trail Name	A two character trail name.	All
Trail Subdirectory	Where you want the trail information to be stored or the name of the log file. The default trail file location is the Oracle GoldenGate Data Home.	All
Begin	How you want the Replicat to start. At a custom time that you select, a specific log position, or the Now default.	All
Transaction Log Sequence Number	Set the log file number of the log file that contains the `START TRANSACTION` record for the transaction you are creating.	All
Transaction Log RBA Offset	Set the record offset of the log file you specified.	All
Checkpoint Table	Set the use of an existing checkpoint table.	All
Max Threads Number	Set a hard limit for the maximum number of active threads that can run simultaneously or use the 25 threads default.	Coordinated

4. In the Database Credentials section, specify the database credentials for applying the trail data. If the database credentials don't exist, then create the credentials for the database instance for Replicat.
5. In the Encryption Profile section, enter the following details:

Options	Description
Profile Name	Name of the profile associated with the Replicat process being created. By default, the LocalWallet profile is selected. If you want to select a custom profile, then select it from the drop-down list.
Encryption Profile Type	This option cannot be changed for the local wallet.
Masterkey Name	This is the default master key for the local wallet. You cannot edit this value.

6. In the Managed Processes section, the options to enter are:

Option	Description
Profile Name	Provides the name of the autostart and autorestart profile. You can select the default or custom options. If you have already created a profile, then you can select that profile also. If you select the Custom option, then you can set up a new profile from this section itself.
Critical to deployment health	(Oracle only) Enable this option if the profile is critical for the deployment health. **Note:** This option only appears while creating the Extract or Replicat and not when you set up the managed processes in the Profiles page.
Auto Start	Enables autostart for the process.
Startup Delay	Time to wait in seconds before starting the process
Auto Restart	Configures how to restart the process if it terminates
Max Retries	Specify the maximum number of retries to try to start the process
Retry Delay	Delay time in trying to start the process
Retries Window	The duration interval to try to start the process
Restart on Failure only	If true the task is only restarted if it failes
Disable Task After Retries Exhausted	If true then the task is disabled after exhausting all attempts to restart the process.

Click **Next**.

7. Enter the Replicat parameters on the Replicat Parameters page. See the Quickstart: Set Up Data Replication with Oracle GoldenGate Microservices Architecture to know more.

8. Click **Create and Run** to start the Replicat process. A green check mark will appear next to the Replicat on the Administration Service Overview page when the Replicat starts successfully.

   You can also check the **Notification** bell icon on the top-right corner of the page, to know the status of Replicat or any other processes.

- Basic Parameters for Parallel Replicat
- Using Replicat Actions

4.4.2.1 Basic Parameters for Parallel Replicat

The following table lists the basic parallel Replicat parameters and their description.

Parameter	Description
MAP_PARALLELISM	Configures number of mappers. This controls the number of threads used to read the trail file. The minimum value is 1, maximum value is 100 and the default value is 2.
APPLY_PARALLELISM	Configures number of appliers. This controls the number of connections in the target database used to apply the changes. The default value is 4.
MIN_APPLY_PARALLELISM MAX_APPLY_PARALLELISM	The Apply parallelism is auto-tuned. You can set a minimum and maximum value to define the ranges in which the Replicat automatically adjusts its parallelism. There are no defaults. Do *not* use with APPLY_PARALLELISM at the same time.
SPLIT_TRANS_REC	Specifies that large transactions should be broken into pieces of specified size and applied in parallel. Dependencies between pieces are still honored. Disabled by default.
COMMIT_SERIALIZATION	Enables commit FULL serialization mode, which forces transactions to be committed in trail order.
Advanced Parameters	
LOOK_AHEAD_TRANSACTIONS	Controls how far ahead the Scheduler looks when batching transactions. The default value is 10000.
CHUNK_SIZE	Controls how large a transaction must be for parallel Replicat to consider it as large. When parallel Replicat encounters a transaction larger than this size, it will serialize it, resulting in decreased performance. However, increasing this value will also increase the amount of memory consumed by parallel Replicat.

Example Parameter File

```
REPLICAT repe
USERIDALIAS ggwest DOMAIN OracleGoldenGate
```

```
MAP_PARALLELISM 2
MIN_APPLY_PARALLELISM 2
MAX_APPLY_PARALLELISM 10
SPLIT_TRANS_RECS 10.000
MAP *.*, TARGET *.*;
```

4.4.2.2 Using Replicat Actions

You can manage a Replicat process using the **Action**, **Details** option of the Replicat from the Administration Service Overview page.

You can change the status of the Replicat process using the **Action** button to start or stop a Replicat or manage the Replicat process from the **Details** option:

Action	Result
Details	Displays the Process Information page that has the following details: • **Process Information** • **Checkpoint** • **Statistics** • **Parameters** • **Report**
Start/Stop	The Replicat starts or stops immediately.
Start/Stop (in the background)	The Replicat is started or stopped using a background process.
Start with Options	Allows you to change the Replicat start point, CSN, filter duplicates, and threads options, then starts the Replicat.
Force Stop	The Replicat is immediately, forcibly stopped.
Alter	Allows you to change when the Replicat begins, the description, and the intent. It does not start the Replicat.
Delete	Allows you to Start/Stop all Replicats at the same time in the background, if there are more than one Replicat processes.

When you change the status, the list options change accordingly. As status are changing, the icons change to indicate the current and final status. The events are added to the Critical Events table. Additionally, progress pop-up messages appear in the bottom of your browser.

Reviewing the Process Information

Process Information
Displays Replicat process details such as status of Replicat as running or stopped. You can also edit the encryption profile and managed options for auto start and auto restart from here.

Checkpoint
Displays the checkpoint log name, path, timestamp, sequence, and offset value. You can click the Checkpoint Detail icon to view elaborate information about the checkpoint.

Statistics
Displays the active replication maps along with replication statistics based on the type of Replicat.

Parameters
Displays the parameters configured when the Replicat was added. You can change these parameters to adjust your Replicat.

Report
Displays the details about the Replicat including the parameters with which the replicat is running, and run time messages.

4.5 Add and Manage Profiles

Oracle GoldenGate Administration Service provides options to set up profiles for managed processes (Extract and Replicat) and for Key Management Systems:

- The profile for Managed Processes is related to **Autostart** and **Autorestart** properties to control the Extract and Replicat process life cycles. To set managed processes, see Configure Managed Processes.
- The profile for Key Management Systems are related to Oracle Key Vault (OKV) and Oracle Cloud InfraStructure (OCI). There is also the option to manage a local wallet. For steps to create and manage your choice of KMS, see Configure an Encryption Profile.

Topics:

- Configure Managed Processes
- Configure an Encryption Profile

4.5.1 Configure Managed Processes

Oracle GoldenGate Administration Service provides options to set up managed Extract and Replicat (ER) processes. These processes are assigned auto-start and auto-restart properties to control their life cycles.

You can create profiles for managed processes using the Administration Service or the Admin Client. To create a profile in the Administration Service, perform the following tasks:

1. Click Profile from the Administration Service navigation pane.
2. In the Managed Process Settings tab, you can click + sign to start creating a profile. There's also a default profile preset on this page.
3. Enter the details for the profile options including the Profile Name, Description, Auto Start and Auto Restart options. See the following table for Auto Start and Auto Restart options.

Option	Description
Profile Name	Provides the name of the autostart and autorestart profile. You can select the default or custom options. If you have already created a profile, then you can select that profile also. If you select the Custom option, then you can set up a new profile from this section itself.
Critical to deployment health	(Oracle only) Enable this option if the profile is critical for the deployment health. **Note:** This option only appears while creating the Extract or Replicat and not when you set up the managed processes in the Profiles page.
Auto Start	Enables autostart for the process.
Startup Delay	Time to wait in seconds before starting the process
Auto Restart	Configures how to restart the process if it terminates
Max Retries	Specify the maximum number of retries to try to start the process
Retry Delay	Delay time in trying to start the process
Retries Window	The duration interval to try to start the process
Restart on Failure only	If true the task is only restarted if it failes
Disable Task After Retries Exhausted	If true then the task is disabled after exhausting all attempts to restart the process.

4.5.2 Configure an Encryption Profile

Oracle GoldenGate Administration Service provides options to set up encryption profiles for managed Extract and Replicat (ER) processes.

To set up the encryption profile, click Profile from the navigation pane and then select the Key Management System (KMS) tab.

1. By default, the Local Wallet profile is created. If you select the Local Wallet encryption profile, you'll see its options, which you can edit using the pen icon.

Options	Description
Description	A description of the local wallet.
Default Profile	This option is enabled by default. You can select to disable it.
Encryption Profile Type	This option cannot be changed for the local wallet.

Options	Description
Masterkey Name	This is the default master key for the local wallet. You cannot edit this value.
Masterkey Version	This is the master key version number. The value is set to **LATEST** and cannot be changed.

2. Click the + sign next to Profile to create an encryption profile by specifying the following details:

Option	Description
Profile Name	Name of the encryption profile
Description	Describe the encryption profile.
Default Profile	If you want to make this profile the default, then enable this option.
Encryption Profile Type	Available options are Oracle Key Vault (OKV) and Oracle Cloud Infrastructure (OCI).

3. Before you set up OKV, you need to perform a client installation. See Step 1: Configure the Oracle Key Vault Server Environment in the *Oracle Key Vault Administrator's Guide*.

OKV Configuration Options	**Options** to set up Oracle Key Vault (OKV)
KMS Library Path	Specify the directory location where Oracle Key Vault is installed.
Oracle Key Vault Version	Specify the supported Oracle Key Vault version.
Masterkey Name	Specify the name of the master key
Time to Live	Time to live (TTL) for the key retrieved by Extract from KMS. When encrypting the next trail, Extract checks if TTL has expired. If so, it retrieves the latest version of the master key. The default is 24 hours.

4. OCI KMS requires registering of the private/public key (API Signing Key) for accessing the REST API on the Server on which Oracle GoldenGate is deployed. Here are the options to set up the OCI KMS encryption profile. See Registering and Managing Keys for OCI KMS.

OCI KMS Configuration Options	**Options** to set up an OCI KMS.
Crypto Endpoint URL	You can access this from the OCI KMS Vault wizard. See OCI Command Line Reference and Managing Keys in *OCI Documentation* to know more.
Tenancy OCID	When you sign up for Oracle Cloud Infrastructure, Oracle creates a *tenancy* for your company, which is a secure and isolated partition within Oracle Cloud Infrastructure where you can create, organize, and administer your cloud resources. See Key Concepts in *OCI Documentation* to learn more.

Key OCID	See the OCI Documentation for details.
User OCID	See the OCI Documentation for details.
API Key	A credential for securing requests to the Oracle Cloud Infrastructure REST API.
API Key Fingerprint	See Required Keys and OCIDs in the OCI documentation for details.

4.6 Access Extract and Replicat Log Information

The diagnosis of Extract and Replicat transactions provides information about the severity of a transaction along with the timestamp. This information is helpful in case you need to determine if and when a particular issue occurred including the cause of the issue.

The Extract and Replicat log information is available on the Diagnosis page of Administration Service. To access the Diagnosis page, click the **left navigation page** of the Administration Service and select **Diagnosis**.

Using the Table

An updated log of connected distribution path and target initiated paths is displayed. You can sort the list by date or severity by clicking on the adjacent arrow. Also, you can refresh this log and choose how many pages you want to view.

To search, you select Date, Severity, or Message, and then select the appropriate options to construct your search.

Notice the **Notifications** tab at the bottom of the page. It displays messages from the service, which are not updated in the log due to transaction errors. For example, failure to log in to the database using the database credentials.

4.7 Enable and Use Debug Logging

You can enable debug logging and download debug log files from this page.

Enabling Debug Logging:

To enable debug logging:

1. Click the Debug Log option from the Navigation Pane of the Service Manager page.
2. Click the Enable Debug Log option to start logging debug information.

Using the Debug Log

You can use the access and use the debug log file from this page:

1. Click the **Download Log File** option to save a local copy of the debug log
2. Click the **Load Debug Log File** option to view the debug log on this page.
3. Search for specific entries in the debug log using the **Search By** box, if required. You can click **Refresh** to get the latest log information, if it doesn't get refreshed automatically.

.

4.8 Add Oracle GoldenGate Deployment Users from the Administration Service

Oracle GoldenGate MA users can be created from the Administration Service, once you log in using the credentials created at the time of configuring the deployment.

To create a user, perform the following tasks:

1. Click **Administrator** from the left navigation pane of the Administration Service.
2. Click **+** to add a user.
3. Enter the required credentials in the fields.
4. Make sure that you select a role from the **Role** drop-down list. The available roles are: Security, Administrator,User, and Operator.
5. Click **Submit**.

   The new user is listed in the Users table including the role and information that you supplied.

5

Configure Paths to Transport Trail Data

This section discusses the steps to create distribution paths (DISTPATHS) and receiver paths (RECVPATHS) to transport trail data from the Distribution Service to the Receiver Service or from the Receiver Service to the Distribution Service.

Topics:

- Quick Tour of the Distribution Service Overview Page
- Add and Manage a Distribution Path
- Add a Target-Initiated Distribution Path
- Review the Distribution Path Information

5.1 Quick Tour of the Distribution Service Overview Page

The Distribution Service is accessible from the Service Manager home page or by directly specifying the URL.

Log in to the Distribution Service with the available user credentials. From the Distribution Service Overview page you can see dashboard that dislays the path that connects the Extract and Replicat processes. You can also add paths from this page. Use the dashboard to perform the following opeartions.

Action	Task
Add paths	See Adding New Paths
View path details	See Using the Path Actions
Start or Stop the path	See Using the Path Actions
Reposition the path	See Using the Path Actions
Enable sharding using filters	See Using the Path Actions and also Adding New Paths
Set or customize the DML filtering	See Using the Path Actions and also Adding New Paths
Set the DDL filtering	See Using the Path Actions and also Adding New Paths
Set or customize Procedure filtering	See Using the Path Actions and also Adding New Paths
Customize Tag filtering	See Adding New Paths
Delete a Path	See Using Path Actions

- Reposition a Path
- Change the Path Filtering

5.1.1 Reposition a Path

You can reposition a path whenever it’s necessary.

On the Distribution Service Overview page, click the Action button for the path. From the drop-down list, click **Reposition**.

Change one or both of the source database options to reposition the path, then apply the changes.

5.1.2 Change the Path Filtering

If you want to change the filter settings for an existing path, the steps are mostly the same as those for creating the filtering for a new path.

On the Distribution Service Overview page, click Action button for the path. From the drop-down list, click **Change Filtering**.

Rule Configuration	Description
Enable filtering	If you enable filtering by selecting it from the toggle button and click the `Add Rule` button, you’ll see the Rule Definition dialog box. • `Rule Name` • `Rule Action`: Select either Exclude or Include • `Filter Type`: Select from the following list of options: – `Object Type`: Select from three object types: `DML`, `DDL`, and `Procedure` – `Object Names:` Select this option to provide an existing object name. A 3–part naming convention depends on whether you are using CDB. With CDB, you need to use a 3–part naming convention, otherwise a 2–part convention is mandatory. 3–part convention includes container, *`schema`*, *`object`*. 2–part convention includes *`schema`*, *`object name`*. – `Procedure Feature Name`: Select this option to filter, based on existing procedure feature name. – `Column Based`: If you select this option, you are presented with the option to enter the table and column name to which the rule applies. You can filter out using column value with `LT`, `GT`, `EQ`, `LE`, `GE`, `NE` conditions. You can also specify if you want to have before image or after image in filtered data. – `Tag:` Select this option to set the filter based on tags. – `Chunk ID`: Displays the configuration details of database shards, however, the details can’t be edited. • `Negate`: Select this check box if you need to negate any existing rule. You can also see the JSON script for the rule by clicking the JSON tab.

After you add a rule, it is listed in Inclusion Rules. You can delete rules or edit them. When you edit a rule, you have the same options as adding a rule with the following added filters:

Option	Description
OR AND	Select one logical operator.
Chunk ID	Edit or delete the database shard settings if sharding is used.
Object Type:	Edit or delete the type of object for the rule.

5.2 Add and Manage a Distribution Path

Distribution path defines the route for the trail to send and receive data. Paths can be initiated from the Distribution Server to the Receiver Service. Alternatively, in cases where there are network security policies that prevent the Distribution Service to open a network connection in the target endpoint to the Receiver Service, the path is initiated from the Receiver Service to the Distribution Service. These types of paths are called target-initiated paths, which are suitable in environments such as Demilitarized Zone Paths (DMZ) or Cloud to on-premise networks.

You need to choose the type of path you need to set up at the time of adding the deployment to the Service Manager, and use the OGGCA wizard to set up the security options that allow you to add target-initiated paths for the deployment. See step 9 from Add a Deployment .

Topics:

- Add a Distribution Path

5.2.1 Add a Distribution Path

A path is created to send the transaction of data from the Extract to the Replicat. You can create a path from the Distribution Service. To add a path for the source deployment:

1. Log in to the **Distribution Service**.
2. Click the plus (+) sign next to Path on the Distribution Service home page.

   The Add Path page is displayed.
3. Enter the details as follows:

Options	Description
Path Name	Select a name for the path.
Description	Provide a description. For example, the name of the Extract and Replicat processes associated with the distribution path.
Reverse proxy enabled?	Select to use reverse proxy. To know more about configuring you reverser proxy servers, see Reverse Proxy Support in the *Step by Step Data Replication Using Oracle GoldenGate Microservices*.
Source: *Trail Name*	Select the Extract name from the drop-down list, which populates the trail name automatically. If it doesn’t, enter the trail name that you provided while adding the Extract.

Options	Description
Generated Source URI:	A URI is automatically generated for the trail based on the Extract information you provided. You can edit this URI by clicking the pencil, then modifying the source. Typically, you will need to edit the URI if you want to use reverse proxy.
Target Authentication Method	Select the authentication method for the target URI. Authentication options are: **OAuth**: Use the **OAuth** if the source and target deployments are IDCS-enabled. This option uses the client credentials for authentication from the Distribution Service to the Receiver Service. **Certificate**: Choose a certificate from the drop-down. This certificate is created using the Certificate Management page in Service. Manager. See Add and Manage Certificates for a Deployment. **UserID Alias**.
Target	Enter the target endpoint of the path. From the drop-down list, select your data transfer protocol. The default option is **wss** (secure web socket). Specify the following details when you select this option: • Target Host: Enter the URL of the target host, for example, localhost, if the target is on the same system. • Port Number: You may enter the port number of the Receiver Service and the trail name of the Replicat you created earlier. However, it’s not mandatory. The port is the Manager port number for Classic Architecture. • Trail Name: Path takes the source trail and sends the date to a target trail given here, which can be consumed by any Replicats created later. • Domain: Name of the target domain. • Alias: User alias of the target domain. You can also choose **ogg** or **ws** (web socket) protocol. For the **ogg** protocol, you need to specify only the target host, port number, and trail file name. For the **ws** protocol, the options are the same as the wss protocol.
Generated Target URI	A target URI is automatically generated for the trail based on the target authentication method and target you provided. You can edit this URI by clicking the pencil, then modifying the target.

Options	Description
Target Encryption Algorithm	Select the encryption algorithm for the target trail. Options include NONE, AES128, AES192, AES256.
Enable Network Compression	Set the compression threshold value if you enable this option.
Compression Threshold	Option appears when you enable the network compression. Specify the compresion threshold value.
Sequence Length	The length of the trail sequence number.
Trail Size (MB)	The maximum size of a file in a trail.
Configure Trail Format	Toggle this switch to enable and configure the trail file format.
Type	Select one of these types of trail file formats: • Plain Text • XML • SQL
Compatible With	Select the utility that is compatible with the trail file. Options are: • BCP • SQLLOADER • COMCAST
Timestamp Precision	Specify the timestamp precision value for the trail file.
Extra Columns	Includes placeholders for additional columns at the end of each record. Use this option when a target table has more columns than the source table. Specify a value between 1 and 9.
Include SYSKEY	Select this option incase your Replicat configuration includes tables with `SYSKEY`.
Quote Style	Select the quote style depending on the database requirements.
Include Column Name?	Enable this option to include column names in the trail file.
Null Is Space?	Select this option to indicate that any null values in the trail file is a space.
Include Place Holder?	Outputs a placeholder for missing columns.
Include Header Fields?	Select to include header fields in the trail file.
Delimiter	An alternative delimiter character.
Use Qualified Name?	Select to use the fully qualified name of the parameter file.
Include Transaction Info?	Enable to to include transaction information.
Encryption Profile	Section

Options	Description
Begin	Select the point from where you need to log data. You can select the following options from the drop-down list: • Now • Custom Time • Position is Log (default)
Source Sequence Number	Select the sequence number of the trail from source deployment Extract.
Source RBA Offset	This setting provides the Relative Byte Address (RBA) offset value which is the point in the trail file (in bytes) from where you want the process to start.
Critical	The default value is false. If set to true, this indicates that the distribution path is critical to the deployment.
Auto Restart	The default value is false. If set to true, the distribution path restarts automatically if it's terminated.
Auto Restart Options	Section
Retries	The number of times to try an restart the task (path process).
Delay	The duration interval to wait between retries.

Rule Configuration	Description
Enable filtering	If you enable filtering by selecting it from the toggle button and click the `Add Rule` button, you'll see the Rule Definition dialog box. • `Rule Name` • `Rule Action`: Select either Exclude or Include • `Filter Type`: Select from the following list of options: – `Object Type`: Select from three object types: `DML`, `DDL`, and `Procedure` – `Object Names:` Select this option to provide an existing object name. A 3–part naming convention depends on whether you are using CDB. With CDB, you need to use a 3–part naming convention, otherwise a 2–part convention is mandatory. 3–part convention includes container, *`schema`*, *`object`*. 2–part convention includes *`schema`*, *`object name`*. – `Procedure Feature Name`: Select this option to filter, based on existing procedure feature name. – `Column Based`: If you select this option, you are presented with the option to enter the table and column name to which the rule applies. You can filter out using column value with `LT`, `GT`, `EQ`, `LE`, `GE`, `NE` conditions. You can also specify if you want to have before image or after image in filtered data. – `Tag:` Select this option to set the filter based on tags. – `Chunk ID`: Displays the configuration details of database shards, however, the details can't be edited. • `Negate`: Select this check box if you need to negate any existing rule. You can also see the JSON script for the rule by clicking the JSON tab.

Additional Options	Description
Eof Delay (centiseconds)	You can specify the Eof Delay in centiseconds. On Linux platforms, the default settings can be retained. However, on non-Linux platforms, you may need to adjust this setting for high bandwidth, high latency networks, or for networks that have Quality of Service (QoS) settings (DSCP and Time of Service (ToS)).
Checkpoint Frequency	Frequency of the path that is taking the checkpoint (in seconds).
TCP Flush Bytes	Enter the TCP flush size in bytes.
TCP Flush Seconds	Enter the TCP flush interval in seconds.
TCP Options	Section
DSCP	Select the Differentiated Services Code Point (DSCP) value from the drop-down list, or search for it from the list.
TOS	Select the Type of service (TOS) value from the drop-down list.
TCP_NODELAY	Enable this option to prevent delay when using the Nagle’s option.
Quick ACK	Enable this option to send quick acknowledgment after receiving data.
TCP_CORK	Enable this option to allow using the Nagle’s algorithm cork option.
System Send Buffer Size	You can set the value for the send buffer size for flow control.
System Receive Buffer Size	You can set the value for the receive buffer size for flow control.
Keep Alive	Timeout for keep-alive.

4. Click **Create Path** or **Create and Run**, as required. Select **Cancel** if you need to get out of the Add Path page without adding a path.

Once the path is created, you’ll be able to see the new path in the Overview page of the Distribution Service.

- Using the Path Actions

5.2.1.1 Using the Path Actions

After a new path is added, you can perform actions such as stop or pause a path, view reports and statistics, reposition the path, change its filtering, and delete a path.

On the Overview page of the Distribution Service, click the **Action** button next the Distribution Path name. From the drop-down list, you can use the following path actions:

- **Details**: Use this option to view details of the path. You can view the path information including the source and target. You can also edit the description of the path. Statistical data is also displayed including LCR Read from Trails, LCR Sent, LCR Filtered, DDL, Procedure, DML inserts, updates, and deletes, and so on. You can also update the App Options and TCP Options.

- **Start** or **Stop**: Use these options to start or stop a path. If the path isn’t started, the **Start** option is displayed instead of the **Stop** option. For a target-initiated distribution path, you can only stop this path from the Distribution Service and cannot delete or start it from the Distribution Service. After you stop the path, it'll not be available on the Distrbution Service.
- **Delete**: Use this option to delete a path. This option is available only when the path is in stopped state. Click **Yes** on the confirmation screen to complete path deletion.
- **Reposition**: Use this option to change the Source Sequence Number and Source RBA Offset
- **Change Filtering**: Use this option to enter sharding, DML filtering, DDL filtering, Procedure filtering, and Tag filtering options.

Depending on the action you select, you can see the change in status at the bottom of the Overview page.

5.3 Add a Target-Initiated Distribution Path

A target-initiated distribution path is created from the Receiver Service. These paths can be used when communication must be initiated from the target. For an overview of target-initiated distribution path, see Overview of Target-Initiated Paths.

To create a target-initiated distribution path, perform the following steps:

1. Log in to the Receiver Service.
2. Click the + sign on the home page to start adding a path.
3. The following table lists the options to set up the path:

Table 5-1 Add a Target-Initiated Distribution Path

Options	Description
Path Name	Name of the target-initiated distribution path
Description	Provide a description of the path.
Reverse Proxy Enabled	Select to use reverse proxy. To know more about configuring your reverse proxy servers, see Reverse Proxy Support.
Source Authentication Method	Select the authentication method for the source URI. Authentication options are OAuth 2.0, Certificate, UserID Alias.

Table 5-1 (Cont.) Add a Target-Initiated Distribution Path

Options	Description
Source	From the drop-down list, select your data transfer protocol. The default option is Secure Web Socket Proctocl (**wss**). Other option is **ws**. You also need to enter the following details: • Source Host: URL of the source host for example, **localhost**, if the source is on the same system. • Port Number: Enter the port number of the Distribution Service. • Trail Name: Enter the trail name you want to read on your source. **NOTE:** The Distribution Service doesn't not create any trail on source. It can only read the provided trail name. • Domain: Enter the domain for the host. • Alias: Provide an alias for this host. Path takes the source trail and sends the data to a target trail given here, which can be consumed by any Replicats created later.
Generated Source URI	A URI is automatically generated for the trail based on the source information you provided.
Target	Name of the target trail of the Replicat you created earlier.
Generated Target URI	A Target URI is automatically generated for the trail based on target trail information you provided.
Target Encryption Algorithm	Select the encryption algorithm for the target trail. Options include AES128, AES192, AES256.
Enable Network Compression	Set the compression threshold value if you decide enable this option.
Sequence Length	The length of the trail sequence number.
Trail Size	The maximum size of a file in a trail.
Configure Trail Format	Toggle this switch to enable and configure the trail file format.
Type	Select one of these types of trail file formats: • Plain Text • XML • SQL
Compatible With	Select the utility that is compatible with the trail file. Options are: • BCP • SQLLOADER • COMCAST
Timestamp Precision	Specify the timestamp precision value for the trail file.

Table 5-1 (Cont.) Add a Target-Initiated Distribution Path

Options	Description
Extra Columns	Includes placeholders for additional columns at the end of each record. Use this option when a target table has more columns than the source table. Specify a value between 1 and 9.
Include SYSKEY	Select this option incase your Replicat configuration includes tables with `SYSKEY`.
Quote Style	Select the quote style depending on the database requirements.
Include Column Name?	Enable this option to include column names in the trail file.
Null Is Space?	Select this option to indicate that any null values in the trail file is a space.
Include Place Holder?	Outputs a placeholder for missing columns.
Include Header Fields?	Select to include header fields in the trail file.
Delimiter	An alternative delimiter character.
Use Qualified Name?	Select to use the fully qualified name of the parameter file.
Include Transaction Info?	Enable to to include transaction information.
Encryption Profile	Section
Begin	Select the point from where you need to log data. You can select the following options from the drop-down list: • Now • Custom Time • Position is Log (default)
Source Sequence Number	Select the sequence number of the trail from source deployment Extract.
Source RBA Offset	This setting provides the Relative Byte Address (RBA) offset value which is the point in the trail file (in bytes) from where you want the process to start.
Critical	The default value is false. If set to true, this indicates that the distribution path is critical to the deployment.
Auto Restart	The default value is false. If set to true, the distribution path is restarted automatically when killed.
Auto Restart Options	Set up the auto restart option in this section.
Retries	The number of times to try an restart the task (path process).
Delay	The duration interval to wait between retries.

Rule Configuration	Description
Enable filtering	If you enable filtering by selecting it from the toggle button and click the `Add Rule` button, you'll see the Rule Definition dialog box. • `Rule Name` • `Rule Action`: Select either Exclude or Include • `Filter Type`: Select from the following list of options: – `Object Type`: Select from three object types: `DML`, `DDL`, and `Procedure` – `Object Names:` Select this option to provide an existing object name. A 3–part naming convention depends on whether you are using CDB. With CDB, you need to use a 3–part naming convention, otherwise a 2–part convention is mandatory. 3–part convention includes container, *`schema`*, *`object`*. 2–part convention includes *`schema`*, *`object name`*. – `Procedure Feature Name`: Select this option to filter, based on existing procedure feature name. – `Column Based`: If you select this option, you are presented with the option to enter the table and column name to which the rule applies. You can filter out using column value with `LT`, `GT`, `EQ`, `LE`, `GE`, `NE` conditions. You can also specify if you want to have before image or after image in filtered data. – `Tag:` Select this option to set the filter based on tags. – `Chunk ID`: Displays the configuration details of database shards, however, the details can't be edited. • `Negate`: Select this check box if you need to negate any existing rule. You can also see the JSON script for the rule by clicking the JSON tab.

Additional Options	Description
Eof Delay (centiseconds)	You can specify the Eof Delay in centiseconds. On Linux platforms, the default settings can be retained. However, on non-Linux platforms, you may need to adjust this setting for high bandwidth, high latency networks, or for networks that have Quality of Service (QoS) settings (DSCP and Time of Service (ToS)).
Checkpoint Frequency	Frequency of the path that is taking the checkpoint (in seconds).

Additional Options	Description
TCP Flush Bytes	Enter the TCP flush size in bytes.
TCP Flush Seconds	Enter the TCP flush interval in seconds.
TCP Options	Section
DSCP	Select the Differentiated Services Code Point (DSCP) value from the drop-down list, or search for it from the list.
TOS	Select the Type of service (TOS) value from the drop-down list.
TCP_NODELAY	Enable this option to prevent delay when using the Nagle's option.
Quick ACK	Enable this option to send quick acknowledgment after receiving data.
TCP_CORK	Enable this option to allow using the Nagle's algorithm cork option.
System Send Buffer Size	You can set the value for the send buffer size for flow control.
System Receive Buffer Size	You can set the value for the receive buffer size for flow control.
Keep Alive	Timeout for keep-alive.

For target-initiated distribution paths, the use case for the **ws** and **wss** protocols is explained in the following table:

Deployment Type	Target Deployment (Non-Secure)	Target Deployment (Secure)
Source Deployment (Non-secure)	`ws`	`ws`
Source Deployment (Secure)	`wss`	`wss`

The `wss` protocol must be specified whenever the source deployment (Distribution Service host) has been configured with security enabled. The secured communication channel can be created using an SSL certificate in a client Wallet, even if the target deployment (Receiver Service host) has disabled security.

Features and Limitations for Using Target-initiated Distrbution Paths

Here are the limitations when working with target-initiated distribution paths:

- There is no support for interaction between legacy and secure deployments using this mode of operation for target-initiated distribution paths.
- No support for `ogg` protocol. Only `ws` and `wss` protocols are supported.
- It is possible to only get information and stop a target-initiated distribution path on Distribution Service and after the path stops, it is not be visible on the Distribution Service.

You can also set up target-initiated distribution paths using the Admin Client. For command options, see the Admin Client commands `ADD RECVPATH`, `ALTER RECVPATH`, `INFO RECVPATH`, `DELETE RECVPATH`, `START RECVPATH` in Admin Client Command Line Interface Commands.

5.4 Review the Distribution Path Information

You can constantly monitor the activity of the path on the Distribution Service Process Information page.

- The path details that you configured. You can change the Description, source and target URIs, Target Authentication Method, DB Name, Target Encryption Algorithm, Enable Network Compression, Sequence Length, Trail Size, configure trail format, mark as Critical and enable Auto Restart. When changing the trail format, be sure to apply your changes.
- The advanced options are the delay, flush, and TCP that you configured. You can change any or all of these options, then apply to the path.

The Statistics tab shows you detailed information about the path, such as the different path types and tables. You can use the arrows to sort the tables and the search to quickly locate a specific table. The search is case insensitive and starts searching as you type to update the table.

6

Monitor Paths and Trails from the Receiver Service

You can use trails to monitor path, tune networks, and data input and output from the Receiver Service.

This section describes the tasks to monitor trails, tune network parameters, and view statistical data about trails:

Topics:

6.1 Quick Tour of the Receiver Service Overview Page

The Receiver Service is the central control service that handles all incoming trail files.

The Receiver Service works with the Distribution Service to receive incoming trail file information. From the Receiver Service Overview page, you can see the status of the distribution path with one end depicting Extract and the other end, Replicat.

You can use the Receiver Service Overview page to view the path details by clicking the **Action**, **Details** option.

Also see Monitoring Paths.

6.2 Tune the Network Parameters

The network settings in Receiver Service are for target-initiated paths and must mirror the ones in Distribution Service. Network parameters include TCP flush byte options, DSCP, ToS, buffer size settings and other settings

You can monitor and fine-tune these parameters depending on your requirements using the Performance Metrics and Distribution Service. However, this applies to Distribution Service if the path is initiated from the Distrbution Service and to Receiver Service when the path is initiated from the Receiver Service. You can view the network parameters from the Performance Monitor Service Overview page for paths that are initiated from the Distribution Service. If you need to tweak them, go to the Distribution Service and do the following:

1. Click the path **Action**, **Details**.

   The Path Information page is displayed.

2. Expand the Advanced Options.

You’ll see App Options, which contain the TCP Flush Bytes and TCP Flush Seconds values. By default, this value is set to OS Default.

The TCP Options, include the following parameters:

- DSCP
- TOS
- Nodelay
- Quick ack
- Cork
- System Send Buffer Size
- System Receiver Buffer Size

3. Click the **Edit** icon next to **Advanced Options**, to change any of the these values,.
4. Click **Apply** to save the changes to the network parameters.

After you edit the network parameters, monitor their status changes and messages from the service. You can do so using the Performance Monitor Service. See Monitoring Performance for details.
For paths initiated from the Receiver Service, the network statistics can be tweaked from the Receiver Service by performing the following steps:

1. Click the target-initiated path **Action** button and select **Details**.
2. From the Path Information tab, expand the Advanced Options, which has the setting for EoF Delay (centiseconds). You may typically need to edit this setting for non-Linux platfoms.

6.3 Review Statistical Information About Paths

You can constantly monitor the activity of the path on the Receiver Service Statistics page.

The Statistics tab shows you detailed information about the logical change records (LCRs) and DDLs that were read from trails, LCRs and DDLs sent and received, LCRs and DDLs filtered. It also provides information about the DML types, inserts, updates, upserts, and deletes.

The table information includes the values of LCRs read and sent. You can use the arrows to sort the tables and the search to quickly locate a specific table. The search is case insensitive and starts searching as you type to update the table.

6.4 Monitor Paths

You can monitor the path network statistics from the Receiver Service. Use the information provided on this page to troubleshoot performance issues with the Distribution Service. If it's not able to keep up, you can access this page and see the reasons why, and then use that information to tune the TCP window size, enable compression, or split the trails into multiple threads (multiple distribution service paths, each moving a subset of tables).

In the Receiver Service, you’ll see the path depicted in a graphical representation and you can perform the following steps to monitor the selected path:

1. Log in to the **Receiver Service** home page.
2. Click **Action**, **Details** for a running path.
3. Click the Network tab.

   You can review the path statistics from this tab. This page displays the following details:

   - Network Statistics: The network statistics information includes details such as target trail file name, port number, total messages written out, and so on. You can use this information to go back to the Distribution Service and tune the network parameters, if required.
   - File IO Statistics: The file IO statistics include total bytes read and total idle time.

7

Monitor Performance from the Performance Metrics Service

The Performance Metrics Service provides a dashboard view as well as a detailed view of status changes, statistical data of the Services’ performance. They are represented through statistical charts and real-time data.

Topics:

- Quick Tour of the Performance Metrics Service Overview Page
- Monitor Performance Statistics
- Review Status Changes
- Purge the Datastore

7.1 Quick Tour of the Performance Metrics Service Overview Page

The Performance Metrics Service uses the metrics service to collect and store instance deployment performance results. When you arrive at the Performance Metrics Service OVerview page, you see all the Oracle GoldenGate processes in their current state. You can click a process to view its performance metrics. You can also access service messages and status change details from this page.

Here’s a general overview of the tasks that you can perform from this page.

Task	Description
Review Messages	Reviewing Messages from the Messages Overview tab.
Review Status Changes	Click the Review Status Changes tab to review changes in status of a service.

- Review Messages from Messages Overview Tab

7.1.1 Review Messages from Messages Overview Tab

Messages from the Services are displayed in Performance Metrics Service Overview page.

To review the messages sent or received, do the following:

1. From the Service Manager, click **Performance Metrics Service**.

   The Performance Metrics Service Overview page is displayed.

2. Click the **Messages Overview** tab (if it’s not already selected) to see a drill down into all the service messages.

Scroll through the list of messages or search for a specific message by entering the text in the message.

3. Click **Refresh** to get a synchronized real-time list of messages before you start searching. You can also change the page size to view more or fewer messages.

7.2 Monitor Performance Statistics

All the services and processes of the Microservices Architecture can be monitored at drill-down levels to allow trend monitoring and statistical analysis of data. The Performance Metrics Service offers these detailed views with graphical representations of statistical data in real-time.

The Performance Metrics Service Overview page presents a dashboard view of all the Microservices, along with their statuses. If you want to drill down to view the performance of any microservice, click the service name to open the reports page for that microservice.

The page also provides the option to **Pause** or **Clear** the data displayed on the page. To get a snapshot of the trends captured for each of the microservices, see the following table:

Performance Report Tab	Microservice
Process Performance	• Administration Service • Distribution Service • Performance Metrics Service • Receiver Service • Extracts • Replicats
Thread Performance	• Administration Service • Distribution Service • Performance Metrics Service • Receiver Service • Extracts • Replicats
Status and Configuration	• Administration Service • Distribution Service • Performance Metrics Service • Receiver Service • Extracts • Replicats
Service Statistics	• Distribution Service • Performance Metrics Service
Trail Files	• Extracts • Replicats
Database Statistics	• Extracts • Replicats
Procedure Statistics	• Extracts • Replicats
Cache Statistics	Extracts
Queue Statistics	Extracts

Performance Report Tab	Microservice
Bounded Recovery	• Extracts • Replicats

7.3 Review Status Changes

Real-time status changes to microservices can be monitored from the Performance Metrics Service Status Changes Overview tab.

Status change messages show the date, process name, and its status, which could be running, starting, stopped, or killed.

To view status changes, click **Performance Metrics Service** from the Service Manager home page, and then click the **Status Changes Overview** tab. A list of status change messages from the service appears.

If you are searching for specific messages, you can use the search but make sure you click **Refresh** before you search to ensure that you get the updated status for services.

Note that the search messages appear in different colors to differentiate critical and informational messages.

7.4 Purge the Datastore

You can change the datastore retention and purge it from the Performance Metrics Service **Monitoring Commands** tab.

To view status changes, click **Performance Metrics Service** from the Service Manager home page, and then click the **Monitoring Commands** tab.

The current process retention in days displays.

You can enter the number of retention days or use the sliding icon to set the new period from 1 to 365 days, then **Execute** to activate the purge. The details of the purge displays.