

Oracle® Cloud

Oracle Cloud Infrastructure GoldenGate



Latest Release
F61358-01

The Oracle logo, consisting of a solid red square with the word 'ORACLE' in white, uppercase, sans-serif font centered within it.

ORACLE®

Oracle Cloud Oracle Cloud Infrastructure GoldenGate, Latest Release

F61358-01

Copyright © Oracle and/or its affiliates.

Contents

1 Get started

Access Oracle Cloud Infrastructure GoldenGate	1-1
Oracle Cloud Infrastructure GoldenGate Workshops	1-1
Basic taskflows	1-2
Replicate data taskflow	1-2
Data Transforms taskflow	1-3
Stream and analyze taskflow	1-4
Create Oracle Cloud resources	1-5
Create a compartment	1-5
Create a Virtual Cloud Network and subnet	1-6
Create users	1-6
Create groups	1-7
Create policies	1-8
How to create a policy	1-8
Minimum recommended policies	1-9
Quickstarts	1-10
Connect to OCI GoldenGate using a public IP	1-12
Overview	1-12
Before you begin	1-12
Task 1: Create the deployment	1-12
Task 2: Launch the console	1-16
Learn more	1-16
Connect to Oracle Cloud Infrastructure GoldenGate using a private IP	1-16
Overview	1-16
Before you begin	1-16
Option A: Use OCI Bastion	1-17
Option B: Use your own bastion on OCI Compute	1-18
Connect to OCI GoldenGate using a public load balancer	1-19
Overview	1-19
Before you begin	1-19
Task 1: Create a certificate bundle	1-20
Task 2: Create a deployment	1-20
Task 3: Create the load balancer	1-23

Task 4: Create a DNS record	1-24
Task 5: Create OCI Network Security Rules to allow/deny ingress	1-24
Learn more	1-25
Oracle Database quickstarts	1-25
Replicate data between cloud databases in the same region	1-25
Configure bidirectional replication between two cloud databases in the same region	1-33
Replicate data between cloud databases in different regions	1-41
Replicate data between cloud databases in different regions with VCN peering	1-43
Send data from Oracle GoldenGate to OCI GoldenGate	1-49
Send data from OCI GoldenGate to Oracle GoldenGate	1-52
Big Data quickstarts	1-55
Replicate data from Autonomous Database to OCI Object Storage	1-55
Replicate data from Autonomous Database to OCI Streaming	1-61
Replicate data from Autonomous Transaction Processing to Amazon S3	1-66
Replicate data from ATP to Kafka	1-71
Replicate data from Autonomous Transaction Processing to Azure Data Lake Storage	1-75
Replicate data from Autonomous Transaction Process to Azure Synapse	1-80
Replicate data from Autonomous Transaction Processing to Confluent Kafka	1-85
Replicate data from Amazon RDS to OCI Object Storage	1-91
Replicate data from MongoDB to Autonomous JSON Database	1-96
Capture data from Kafka platforms	1-99
Stage and merge data into Autonomous Data Warehouse using OCI GoldenGate	1-101
Replicate data from PostgreSQL to Snowflake	1-106
Send Data from OCI MySQL Heatwave to Azure Event Hubs	1-115
Replicate Data from OCI MySQL Heatwave to Amazon Kinesis	1-119
Replicate Data from OCI MySQL Heatwave to Google Cloud Storage	1-123
Replicate data from PostgreSQL to Google BigQuery	1-127
SQL Server quickstarts	1-133
Replicate data from Azure SQL Managed Instance to Autonomous Transaction Processing	1-133
MySQL quickstarts	1-139
Replicate data from OCI MySQL Heatwave Database to Autonomous Data Warehouse	1-139
Send Data from OCI MySQL Heatwave to Azure Event Hubs	1-146
Replicate Data from OCI MySQL Heatwave to Amazon Kinesis	1-150
Replicate Data from OCI MySQL Heatwave to Google Cloud Storage	1-154
PostgreSQL quickstarts	1-159
Replicate data from PostgreSQL to Autonomous Transaction Processing	1-159
Replicate data from PostgreSQL to MySQL	1-168
Replicate data from PostgreSQL to Snowflake	1-176
Replicate data from PostgreSQL to Google BigQuery	1-185

Stream Analytics quickstarts	1-191
Create a File stream pipeline	1-191
Replicate data into Stream Analytics	1-194

2 Overview

About Oracle Cloud Infrastructure GoldenGate	2-1
OCI GoldenGate concepts	2-1
Data Transforms concepts	2-2
Stream analytics concepts	2-2
User roles	2-2
OCI GoldenGate Connectivity	2-3
Example: Replication from Autonomous Transaction Processing into Autonomous Data Warehouse	2-5
Example: Replication from Oracle on-premise into Autonomous Data Warehouse	2-6
Example: Replication from Azure SQL Managed Instance into Autonomous Transaction Processing	2-7
Shared responsibility model	2-8
What's New in Oracle Cloud Infrastructure GoldenGate	2-9
December 2023	2-10
October 2023	2-10
August 2023	2-12
July 2023	2-12
June 2023	2-13
March 2023	2-13
February 2023	2-14
December 2022	2-14
November 2022	2-14
October 2022	2-15
September 2022	2-15
August 2022	2-15
June 2022	2-15
March 2022	2-16
December 2021	2-16
November 2021	2-16
October 2021	2-17

3 Plan

Before you begin with Oracle Cloud Infrastructure GoldenGate	3-1
Learn Oracle Cloud	3-1
Region availability	3-1

Service Limits	3-1
Compartment Quotas	3-2
What's next?	3-2
What's supported	3-2
Deployment types	3-2
Supported connection types for data replication	3-2
Learn more	3-4
Supported connection types for Data Transforms	3-5
Supported connection types for Stream Analytics	3-5
OCPU management and billing	3-5
Metering and billing	3-6
About OCPU utilization	3-6
Metering and billing for Stream Analytics deployments	3-7
Integrated services	3-8
IAM	3-8
Events	3-8
Work Requests	3-9
Monitoring	3-9
Example OCI GoldenGate topologies	3-9
How many resources do I need?	3-9
Deployments	3-9
Connections	3-10
Example: Azure SQL Managed Instance to Autonomous Transaction Processing	3-10
Example: Autonomous Transaction Processing to Apache Kafka	3-11
Example: PostgreSQL to Autonomous Transaction Processing	3-11
Example: PostgreSQL to MySQL	3-12
Example: PostgreSQL to Snowflake	3-12

4 Connect

Explore connections	4-1
What is a connection?	4-1
Supported connection types for data replication	4-2
Supported connection types for Data Transforms	4-3
Supported connection types for Stream Analytics	4-4
Oracle Database connections	4-4
Connect to Oracle Autonomous Database	4-5
Connect to Autonomous Database Shared	4-5
Known issues	4-6
Creating a connection to Oracle Database	4-8
Create the connection	4-8

Known issues	4-10
Connect to Oracle Exadata	4-11
Connect to Amazon RDS Oracle Database	4-12
Big Data connections	4-14
Create a connection to Oracle Autonomous JSON Database	4-15
Connect to OCI Object Storage	4-16
Create the connection	4-16
Known Issues	4-17
Connect to OCI Streaming	4-17
Before you begin	4-17
Create the connection	4-18
Connect to Oracle NoSQL	4-19
Connect to Oracle Weblogic JMS (Java Message Service)	4-20
Connect to Amazon Kinesis	4-21
Connect to Amazon Redshift	4-22
Connect to Amazon S3	4-24
Connect to Apache Kafka	4-25
Create the connection	4-25
Troubleshoot Kafka connection errors	4-27
Connect to Confluent Kafka	4-27
Create the connection	4-27
Create a connection to Confluent Cloud with Private Links	4-29
Troubleshoot Kafka connection errors	4-30
Connect to Confluent Schema Registry	4-30
Connect to Azure Cosmos DB for MongoDB	4-32
Connect to Azure Event Hubs	4-33
Connect to Azure Data Lake Storage	4-35
Create the Connection	4-35
Troubleshoot connection issues	4-36
Connect to Azure Synapse Analytics	4-37
Connect to MongoDB	4-38
Create the connection	4-38
Known issues	4-39
Connect to Redis	4-39
Connect to Snowflake	4-40
Create a connection to Google Cloud Storage	4-41
Connect to Google BigQuery	4-42
Connect to Elasticsearch Server	4-43
Connect to Hadoop Distributed File System	4-44
MySQL connections	4-45
Connect to OCI MySQL Heatwave	4-45

Connect to MySQL Database Server	4-47
Connect to Amazon RDS for MySQL	4-48
Connect to Amazon RDS for MariaDB	4-49
Connect to Amazon Aurora MySQL	4-50
Connect to Azure Database for MySQL	4-52
Connect to Google Cloud SQL for MySQL	4-53
Connect to MariaDB	4-54
Connect to SingleStoreDB	4-55
Connect to SingleStoreDB Cloud	4-57
PostgreSQL connections	4-58
Create a connection to PostgreSQL Server	4-58
Connect to Amazon RDS PostgreSQL	4-59
Connect to Amazon Aurora PostgreSQL	4-61
Connect to Azure Database for PostgreSQL	4-62
Connect to Google Cloud SQL for PostgreSQL	4-63
SQL Server connections	4-64
Connect to Azure SQL Database	4-64
Create a connection to Azure SQL Managed Instance	4-65
Create the connection	4-65
Known issues	4-67
Connect to Microsoft SQL Server	4-67
Connect to Amazon RDS for SQL Server	4-68
Connect to Google Cloud SQL for SQL Server	4-69
Create a Generic connection	4-71
Connect to Oracle GoldenGate deployments	4-72
About Oracle GoldenGate connections	4-72
Create a connection to GoldenGate	4-72

5 Replicate data

Create data replication resources	5-1
Create a deployment	5-1
Assign a connection to a deployment	5-5
Access the deployment	5-5
Explore the OCI GoldenGate deployment console	5-6
Add Extracts	5-7
RDBMS Extracts	5-7
Add an Extract for Oracle Database	5-7
Add an Extract for Microsoft SQL Server	5-10
Add an Extract for MySQL	5-12
Add an Extract for PostgreSQL	5-14

Big Data Extracts	5-16
Add an Extract for Kafka	5-16
Add an Extract for MongoDB	5-20
Add a Distribution Path	5-21
When to use a Distribution Path	5-21
Before you begin	5-22
Create and run a Distribution Path	5-22
Learn more	5-25
Known issues	5-25
Add a Receiver Path	5-26
When to use target-initiated Receiver Paths	5-26
Before you begin	5-26
Create and run a Receiver Path	5-26
Learn more	5-29
Add Replicats	5-29
RDBMS Replicats	5-30
Add a Replicat for Oracle Database	5-30
Add a Replicat for Microsoft SQL Server	5-32
Add a Replicat for MySQL	5-33
Add a Replicat for PostgreSQL	5-35
Big Data Replicats	5-36
Add a Replicat for Autonomous Database	5-36
Add a Replicat for Oracle Autonomous JSON Database	5-38
Add a Replicat for an Object Storage target	5-39
Add a Replicat for an OCI Streaming target	5-42
Add a Replicat for Amazon Kinesis	5-43
Add a Replicat for Amazon Redshift	5-45
Add a Replicat for Amazon S3	5-46
Add a Replicat for Kafka	5-48
Add a Replicat for Confluent Kafka	5-49
Add a Replicat for Elasticsearch Server	5-50
Add a Replicat for Google BigQuery	5-50
Add a Replicat for Google Cloud Storage	5-52
Add a Replicat for Azure Cosmos DB for MongoDB	5-54
Add a Replicat for Azure Data Lake Storage	5-55
Add a Replicat for Azure Event Hubs	5-57
Add a Replicat for Azure Synapse Analytics	5-57
Add a Replicat for MongoDB	5-59
Add a Replicat for Redis	5-60
Add a Replicat	5-61
Using the Admin Client	5-63

Access Admin Client	5-63
Connect to Admin Client through Cloud Shell	5-63
Use Admin Client	5-65

6 Stream and analyze

About Stream Analytics	6-1
Stream analytics concepts	6-1
Stream analytics limitations	6-1
Supported connections	6-2
Metering and billing for Stream Analytics deployments	6-2
Create Stream Analytics resources	6-3
Create the Stream Analytics deployment	6-4
Supported connections	6-6
Assign a connection to a deployment	6-7
Access the Stream Analytics deployment	6-7
Use Stream Analytics	6-8
Stream and analyze taskflow	6-8

7 Transform data

About Data Transforms	7-1
Data Transforms concepts	7-1
Supported connection types for Data Transforms	7-1
Create Data Transforms resources	7-2
Create a Data Transforms deployment	7-2
Assign a connection to a deployment	7-4
Access the Data Transforms deployment	7-5
Use Data Transforms	7-5
Data Transforms taskflow	7-5

8 Manage

Manage connections	8-1
View connection details	8-1
Edit a connection	8-2
Assign a deployment to a connection	8-3
Unassign a deployment	8-3
Test connections	8-4
Move a connection	8-4
Manage tags for a connection	8-4
Delete a connection	8-4

Known issues	8-5
Managing deployments	8-6
View deployment details	8-6
Edit a deployment	8-8
Edit a deployment username	8-9
Edit a password secret	8-10
Upgrade a deployment	8-10
Set up notifications	8-10
Scale a deployment	8-11
Collect diagnostics	8-12
Stop a deployment	8-13
Start a deployment	8-13
Move a Deployment	8-13
Manage tags for a deployment	8-14
Delete a deployment	8-14
Assign a connection to a deployment	8-14
Test connections	8-15
Unassign a connection	8-15
Manage deployment users	8-15
In IAM-enabled tenancies	8-16
Configure Identity domains for OCI GoldenGate	8-16
In non-IAM enabled tenancies	8-17
Add a user to a deployment	8-17
Edit a deployment user	8-19
Delete a deployment user	8-19
Managing Trail files	8-20
View Trail files	8-20
Purge Trail files	8-20
Manage master encryption key wallets	8-22
Add a master key in the deployment console	8-22
Export a master encryption key wallet from an OCI GoldenGate deployment	8-23
Export a master key encryption wallet from an on premise Oracle GoldenGate instance	8-23
Import a master encryption key wallet to a deployment	8-24
Import a master encryption key wallet to an on premise GoldenGate instance	8-25
Managing Truststore certificates	8-25
What is a Truststore certificate?	8-25
Add a Truststore certificate	8-26
Delete a Truststore certificate	8-26
Manage deployment backups	8-26
Creating a Manual Backup	8-27
Manual backup directory structure	8-27

Viewing deployment backup details	8-28
Copy deployment backup	8-29
Restoring a deployment from a backup	8-30
Creating a Deployment Clone	8-30
Monitor performance	8-30
Monitor performance in the Oracle Cloud console	8-30
Metrics	8-30
Create Alarms	8-32
Subscribe to Events	8-33
Monitor performance using the OCI GoldenGate deployment console	8-33
Monitoring Oracle GoldenGate Service Performance	8-33
Reviewing Messages	8-35
Review Status Changes	8-35

9 Upgrade

Maintain OCI GoldenGate deployments	9-1
About GoldenGate versions	9-1
Deprecation of versions	9-2
Schedule upgrades	9-2
Before you upgrade	9-3
Upgrade a deployment	9-4
Upgrade notifications	9-5
Snooze notifications	9-5
Rollback upgrades	9-6
Reschedule upgrades	9-6
Cancel upgrades	9-7

10 Secure

Securing OCI GoldenGate	10-1
Responsibilities	10-1
Recommendations	10-2
Examples	10-2

11 Troubleshoot

Troubleshoot using the Oracle Cloud console	11-1
Deployment Information	11-1
Metrics	11-1
Example: Troubleshooting deployment health	11-2
Example: Troubleshooting OCPU Utilization	11-3

Troubleshoot using logs	11-3
Process and error logs	11-4
Add policies to use OCI Logging with OCI GoldenGate	11-4
Enable Logging using the Oracle Cloud console	11-4
Enable OCI Logging using CLI	11-5
Enable Logging for OCI GoldenGate in the OCI Logging service	11-9
Troubleshoot health issues using deployment backups	11-10
Collect diagnostics	11-10
Troubleshoot using the OCI GoldenGate deployment console	11-11
Troubleshoot using the Administration Service	11-11
Troubleshoot using the Distribution Service	11-12
Troubleshoot using the Receiver Service	11-12
Troubleshoot using the Performance Metrics Service	11-13
Troubleshoot memory consumption issues	11-13
Troubleshoot connectivity issues with Oracle Database	11-14
Get help	11-15
Submitting a Service Request	11-15

12 Reference

Frequently asked questions	12-1
About Region Availability	12-1
About Host Names and IPs	12-1
About Ports and Protocols	12-2
About Oracle Databases	12-2
About Oracle Autonomous Databases	12-2
About APIs and SDKs	12-2
Oracle Cloud Infrastructure GoldenGate versions	12-2
About GoldenGate versions	12-2
Deprecation of versions	12-3
Upgrade notifications	12-4
Oracle GoldenGate for Oracle versions	12-4
Oracle GoldenGate for Big Data versions	12-7
Oracle GoldenGate for Microsoft SQL Server versions	12-8
Oracle GoldenGate for MySQL versions	12-10
Oracle GoldenGate for PostgreSQL versions	12-12
Oracle Data Transforms versions	12-14
Oracle Cloud Infrastructure GoldenGate Events	12-14
Deployment Event Types	12-14
Database Registration Event Types	12-16

Deployment Backup Event Types	12-17
Deployment Lifecycle Event Types	12-18
Upgrade Notification Event Type	12-19
Oracle Cloud Infrastructure GoldenGate Metrics	12-20
Overview	12-20
Prerequisites	12-20
Available Metrics	12-21
Using the Console	12-25
Oracle Cloud Infrastructure GoldenGate Policies	12-25
Create policies	12-25
How to create a policy	12-26
Minimum recommended policies	12-26
Policy Examples for Securing Network Resources	12-27
Resource-Types	12-28
Supported Variables	12-29
Details for Verbs + Resource-Type Combinations	12-29
goldengate-deployments	12-29
goldengate-connections	12-30
goldengate-connection-assignments	12-30
goldengate-deployment-backups	12-31
Permissions Required for Each API Operation	12-31
Known Issues in Oracle Cloud Infrastructure GoldenGate	12-33
General	12-33
Deployment console	12-34
Connections	12-34
GoldenGate processes	12-37

1

Get started

Learn how to get started with Oracle Cloud Infrastructure GoldenGate, including how and where to access the service, guided workshops and tutorials to lead you through common use cases, and a basic taskflow of operations to complete when using OCI GoldenGate for the first time.

Articles in this section

- [Access Oracle Cloud Infrastructure GoldenGate](#)
- [Oracle Cloud Infrastructure GoldenGate workshops](#)
- [Basic taskflows](#)
- [Create Oracle Cloud resources](#)

Access Oracle Cloud Infrastructure GoldenGate

Learn to access and find OCI GoldenGate in the Oracle Cloud console.

To access OCI GoldenGate:

1. Use a supported browser to access the Oracle Cloud console (<https://cloud.oracle.com>).
2. Click **Sign In using a Cloud Account Name**.
3. Enter your Cloud Account Name (tenancy name), and then click **Next**.
4. Under Oracle Cloud Infrastructure Direct Sign-In, enter your **User Name** and **Password**, and then click **Sign In**.
5. Open the navigation menu, click **Oracle Database**, and then click **GoldenGate**.

Note:

You can also enter GoldenGate in the search bar.

Oracle Cloud Infrastructure GoldenGate Workshops

Learn to use and configure different OCI GoldenGate replication scenarios through Oracle LiveLab workshops.

Oracle LiveLabs lets you try OCI GoldenGate without having your own tenancy. Select a workshop, click **Start**, and then click **Reserve on LiveLabs Sandbox** to get started:

- [Replicate data using OCI GoldenGate](#)
- [Set up bidirectional replication in OCI GoldenGate](#)
- [Send data from Oracle GoldenGate to OCI GoldenGate](#)
- [Send data from OCI GoldenGate to Oracle GoldenGate](#)

- [Replicate data from Autonomous Transaction Processing to OCI Object Storage](#)

You can run all of the above workshops on your own tenancy as well. The following LiveLab is available only to run on your own tenancy:

- [Replicate data from OCI MySQL database to Autonomous Data Warehouse](#)

Basic taskflows

The following taskflows guide you on how to get set up with OCI GoldenGate and Stream Analytics quickly and easily. Depending on your use case, you may diverge from the path outlined below, in which case you can choose to follow one of the various quickstarts available.

Replicate data taskflow

Task	Description	More information
Review Security best practices	Develop a firm understanding of your responsibilities to keep your OCI GoldenGate deployments and connections secure.	Securing OCI GoldenGate
Create OCI resources	This task is typically performed by an administrator. Ensure that the required networking resources are created before you begin.	Create Oracle Cloud resources
Create deployments	A deployment is a container for your OCI GoldenGate resources.	Example OCI GoldenGate topologies Create deployments
Create connections	A connection contains the network connectivity details for a data source or target.	About connections
Assign connections to deployments	To use a connection as a source or target, you must assign it to a deployment.	Create an association between connections and deployments
Launch the OCI GoldenGate deployment console	Create your data replication processes in the OCI GoldenGate deployment console.	Explore the deployment console
Manage deployment users	Add users that can create and manage data replication processes in the deployment console.	Manage deployment users
Create an Extract	An Extract is a process that captures the data from the source that you want to replicate to the target.	Add an Extract
Create a Distribution Path	Use a Distribution Path when you need to replicate data in a distributed deployment environment.	Add a Distribution Path
Create a Replicat	A Replicat is a process that delivers data to the target.	Add a Replicat

Task	Description	More information
Monitor performance	Monitor the overall health of your deployments to ensure your processes run smoothly.	Monitor performance
Maintain trail files	Trail files can add up exponentially over time and take up valuable space. Ensure that you create tasks to purge unused trail files periodically.	Manage Trail files

Data Transforms taskflow

Note:

Data Transforms is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.

Task	Description	More information
Review Security best practices	Develop a firm understanding of your responsibilities to keep your OCI GoldenGate deployments and connections secure.	Securing OCI GoldenGate
Create OCI resources	This task is typically performed by an administrator. Ensure that the required networking resources are created before you begin.	Create Oracle Cloud resources
Create deployments	A deployment is a container for your OCI GoldenGate resources.	Create a Data Transforms deployment
Create Generic connections	Create a generic connection for each Data Transforms data source.	Create a Generic connection
Assign connections to deployments	To use a connection as a source or target, you must assign it to a deployment.	Assign a connection to a deployment
Launch the Data Transforms console	Create your data flows and workflows in Data Transforms.	Access Data Transforms
Create connections	Create connections in Data Transforms to data sources to use in a project.	Work with connections
Create a project	A project is the top-level container, which can include multiple folders to organize your data flows or work flows into logical groups.	Work with Projects

Task	Description	More information
Create and Run a Data Load	A data load allows you to load multiple data entities from a source connection to a target connection.	Create a Data Load Run a Data Load
Monitor Status of Data Loads, Data Flows, and Workflows	When you run a data load, data flow, or workflow Oracle Data Transforms runs jobs in the background to complete the request. You can view the status of the job in the panel on the bottom right of the page.	Monitor Status of Data Loads, Data Flows, and Workflows
Create Data Flows and Workflows	A data flow defines how the data is moved and transformed between different systems. A workflow is made up of multiple flows organized in a sequence in which they must be executed.	Create a Data Flow Create a New Workflow

Stream and analyze taskflow



Note:

Stream Analytics is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.

Task	Description	More information
Review Security best practices	Develop a firm understanding of your responsibilities to keep your OCI GoldenGate deployments and connections secure.	Securing OCI GoldenGate
Create OCI resources	This task is typically performed by an administrator. Ensure that the required networking resources are created before you begin.	Create Oracle Cloud resources
Create deployments	A deployment is a container for your OCI GoldenGate resources.	Example OCI GoldenGate topologies Create deployments
Create connections	A connection contains the network connectivity details for a data source or target.	About connections
Assign connections to deployments	To use a connection as a source or target, you must assign it to a deployment.	Create an association between connections and deployments
Launch the Stream Analytics console	Create your data replication processes in the Stream Analytics console.	Access the deployment

Task	Description	More information
Create a stream	A Stream is a source of continuous and dynamic data. The data can be from a wide variety of data sources such as IoT sensors to information from geospatial services or social networks.	Create Streams
Create a pipeline	A pipeline includes a sequence of data processing stages such as, Query, Pattern, Rule, Query Group, Custom, and Scoring.	Create a Pipeline
Add Business Logic	Transform the input data stream, add business logic to the pipeline to analyze the input data stream.	See Transform See Analyze
Publish the Pipeline	To make the pipeline available for all the users of Oracle Stream Analytics and send data to targets, you must publish a pipeline.	Publishing a Pipeline

Create Oracle Cloud resources

Learn to create a compartment, VCN, subnet, users, and user groups before you get started with Oracle Cloud Infrastructure GoldenGate.

Create a compartment

Compartments let you organize and control access to your cloud resources. It's a logical container that you can use to group related cloud resources together and let specific user groups access.

When you sign up for Oracle Cloud Infrastructure, Oracle creates your *tenancy*, which is the root compartment that holds *all* your cloud resources. You then create additional compartments within your tenancy and the corresponding policies to control access to the resources in each compartment.

To create a compartment:

1. Open the Oracle Cloud console navigation menu, and then click **Identity & Security**.
2. Under **Identity**, click **Compartments**. A list of the compartments you have access to is displayed.
3. Navigate to the compartment where you want to create the new compartment.
 - To create the compartment in the tenancy (root compartment) click **Create Compartment**.
 - To create the compartment in a compartment other than the tenancy (root compartment), click through the hierarchy of compartments until you reach the detail page of the compartment where you want to create the compartment. On the **Compartment Details** page, click **Create Compartment**.
4. In the Create Compartment dialog, complete the fields as follows:

- a. For **Name**, enter a unique name for the compartment, no more than 100 characters (includes letters, numbers, periods, hyphens, and underscores). The name must be unique across all compartments in the tenancy. Avoid entering confidential information.
- b. For **Description**, enter a description that helps distinguish the compartment from others.
- c. For **Parent Compartment**, verify that this is the compartment you where you want to create your compartment. To choose a different compartment, select one from the dropdown.
- d. (Optional) For **Tag Namespace**, you can add a free-form tag to help you search for you resources in the Oracle Cloud console. Click **+ Another Tag** to add more tags.
- e. Click **Create Compartment**.

Your compartment appears in the Compartments list after it's created. You can now create policies and add resources to your compartment.

Create a Virtual Cloud Network and subnet

A virtual cloud network (VCN) is a network that you set up in the Oracle Cloud Infrastructure data centers in a particular region. A subnet is a subdivision of a VCN. OCI GoldenGate requires a VCN and subnet to control traffic to its resources.

To create a VCN and subnet:

1. Open the Oracle Cloud console navigation menu, click **Networking**, and then select **Virtual Cloud Networks**.
2. On the **Virtual Cloud Networking in <compartment-name>** page, click **Start VCN Wizard**.
3. In the Start VCN Wizard dialog, select **VCN with Internet Connectivity**, and then click **Start VCN Wizard**.
4. On the Configuration page, under **Basic Information**, enter a name for **VCN Name**.
5. For **Compartment**, select the compartment where you want to create this VCN.
6. Click **Next**.
7. On the Review and Create page, verify the configuration details, and then click **Create**.

Click **View VCN Details** to verify that both a Public and Private subnet were created.

Create users

Create users to add to groups that can access to your OCI GoldenGate resources.

Before you create users, understand that:

- User names must be unique across all users within your tenancy
- User names are unchangeable
- Users have no permissions until they're placed in a group

To create users:

1. Open the Oracle Cloud console navigation menu, click **Identity & Security**, and then under **Identity**, click **Users**.
2. On the Users page, click **Create User**.
3. On the Create User page, complete the fields as follows:
 - a. For **Name**, enter a unique name or email address for the user.

 **Note:**

The name must be unique across all users in the tenancy. You cannot change this value later. The user name cannot contain spaces, and can only consist of basic Latin letters (ASCII), numerals, hyphens, periods, underscores, +, and @.

- b. For **Description**, enter the user's full name, a nickname, or other descriptive information.
 - c. For **Email**, enter a valid email address for the user for password recovery. This value must also be unique in the tenancy.
4. Click **Create**.

You can then add the user to a group and create policies that give the group access to your resources. For more information about users, see [Managing users](#).

Create groups

A group is a collection of users who require the same type of access to a set of resources or compartments.

Before you create a group, understand that:

- The group name must be unique within the tenancy.
- The group name cannot be changed once created.
- A group has no permissions unless you write at least one permission that gives the group permission to a tenancy or compartment.

To create a group:

1. Open the Oracle Cloud console navigation menu, click **Identity & Security**, and then under **Identity**, click **Groups**.
2. Click **Create Group**.
3. In the Create Group panel:
 - a. For **Name**, enter a unique name for the group.

 **Note:**

Once the group is created, you cannot change the name. The group name must be unique within the tenancy. The group name can be 1 to 100 alphanumeric characters long, upper or lowercase letters, and can contain periods, dashes, hyphens, but no spaces

- b. For **Description**, enter a friendly description.
4. Click **Create Group**.
5. In the Groups list, select the group. You're brought to the group Details page.
6. Click **Add User to Group**.
7. Select a user from the dropdown, and then click **Add User**.

A group doesn't have any permissions until you write a policy that gives the group permission to a compartment or tenancy. For more information about groups, see [Managing groups](#).

Create policies

Policies define what actions members of a group can perform, and in which compartments.

You create policies using the Oracle Cloud console. In the Oracle Cloud console navigation menu, go to **Identity & Security**, and then under **Identity**, and click **Policies**. Policies are written in the following syntax:

```
allow group <identity-domain>/<group-name> to <verb> <resource-type>
in <location> where <condition>
```

- **<identity-domain>**: (Optional) If using OCI IAM for identity management, then include the identity domain of the user group. If omitted, then OCI uses the default domain.
- **<group-name>**: The name of the user group you're giving permissions to
- **<verb>**: Gives the group a certain level of access to a resource-type. As the verbs go from `inspect` to `read` to `use` to `manage`, the level of access increases and the permissions granted are cumulative. To learn about the relationship between permissions and verbs, see [Permissions](#).
- **<resource-type>**: The type of resource you're giving a group permission to work with. There are individual resources, such as `goldengate-deployments` and `goldengate-connections`, and there are resource families, such as `goldengate-family`, which includes both `goldengate-deployments` and `goldengate-connections`. For more information, see [resource-types](#).
- **<location>**: Attaches the policy to a compartment or tenancy. You can specify a single compartment or compartment path by name or OCID, or specify `tenancy` to cover the entire tenancy.
- **<condition>**: Optional. One or more conditions for which this policy will apply.

Learn more about [policy syntax](#).

How to create a policy

To create a policy:

1. In the Console navigation menu, under **Governance and Administration**, go to **Identity**, and then click **Policies**.
2. Click **Create Policy**.

3. Enter a name and description for the policy.
4. In the **Statement** field, enter a policy rule in the following format:

```
allow <subject> to <verb> <resource-type> in <location> where <condition>
```

Conditions are optional. See [Details for Verbs + Resource-Type Combinations](#).

5. (Optional) To add another statement, click **+ Another Statement**.
6. Click **Create**.

For more information about policies, see [how policies work](#), [policy syntax](#), and [policy reference](#).

Minimum recommended policies

At minimum, you need policies to:

- Allow users to *use* or *manage* GoldenGate resources, so that they can work with deployments and connections. For example:

```
allow group <identity-domain>/<group-name> to manage goldengate-family in
compartment <compartment-name>
```

- Allow users to *manage* network resources, so that they can view and select compartments and subnets, and create and delete private endpoints when creating GoldenGate resources. For example:

```
allow group <identity-domain>/<group-name> to manage virtual-network-
family in compartment <compartment-name>
```

Optionally, you can further secure network resources using a combination of granular policies. See [Policy Examples for Securing Network Resources](#).

- Allow users to read the Identity and Access Management (IAM) user and group for validations in IAM enabled tenancies:

```
allow service goldengate to {idcs_user_viewer, domain_resources_viewer}
in tenancy
```

- Oracle Vault, to access customer managed encryption keys. For example:

```
allow group <identity-domain>/<group-name> to manage secret-family in
<location>
allow group <identity-domain>/<group-name> to use keys in <location>
allow group <identity-domain>/<group-name> to use vaults in <location>
allow service goldengate to use keys in <location>
allow service goldengate to use vaults in <location>
```

Depending on whether you intend to use the following services, you may also need to add policies for:

- Oracle Databases, for your source and/or target databases. For example:

```
allow group <identity-domain>/<group-name> to read database-family  
in compartment <compartment-name>
```

```
allow group <identity-domain>/<group-name> to read autonomous-  
database-family in compartment <compartment-name>
```

- Oracle Object Storage, to store manual OCI GoldenGate backups. For example:

```
allow group <identity-domain>/<group-name> to manage objects in  
<location>  
allow group <identity-domain>/<group-name> to inspect buckets in  
<location>
```

- OCI Logging, to access log groups. For example:

```
allow group <identity-domain>/<group-name> to manage log-groups in  
compartment <compartment-name>  
allow group <identity-domain>/<group-name> to manage log-content in  
compartment <compartment-name>
```

The following statement gives a group permission to manage tag-namespaces and tags for workspaces:

```
allow group <identity-domain>/<group-name> to manage tag-namespaces in  
compartment <compartment-name>
```

To add a defined tag, you must have permission to use the tag namespace. To learn more about tagging, see [Resource Tags](#).

For more information and additional example policies, see [OCI GoldenGate Policies](#).

Quickstarts

Quickstarts are guided instructions on how to set up common replication use cases. Some quickstarts have corresponding workshops where you can try out a use case in a sandbox environment before implementing it in your own environment.

General quickstarts

- [Connect to OCI GoldenGate using a public IP](#)
- [Connect to Oracle Cloud Infrastructure GoldenGate using a private IP](#)
- [Connect to Oracle Cloud Infrastructure GoldenGate using a public load balancer](#)

Oracle Database quickstarts

- [Replicate data between cloud databases in the same region](#)
- [Replicate data between cloud databases in different regions](#)
- [Replicate data between cloud databases in different region with VCN peering](#)
- [Configure bidirectional replication](#)

- [Send data from Oracle GoldenGate to OCI GoldenGate](#)
- [Send data from OCI GoldenGate to Oracle GoldenGate](#)

Big Data quickstarts

- [Replicate data from Autonomous Transaction Processing to OCI Object Storage](#)
- [Replicate data from Amazon RDS for Oracle to OCI Object Storage](#)
- [Replicate data from Autonomous Database to OCI Streaming](#)
- [Replicate data from Autonomous Transaction Processing to Apache Kafka](#)
- [Capture data from Kafka platforms](#)
- [Stage and merge data into Autonomous Data Warehouse using OCI GoldenGate](#)
- [Replicate Data from Autonomous Transaction Processing to Confluent Kafka](#)
- [Replicate data from Autonomous Transaction Processing to Azure Data Lake Storage Gen 2](#)
- [Replicate Data from Autonomous Transaction Processing to Azure Synapse Analytics](#)
- [Replicate data from Autonomous Transaction Processing to Amazon S3](#)
- [Replicate data from MongoDB to Autonomous JSON Database](#)
- [Replicate data from PostgreSQL to Snowflake](#)
- [Replicate data from MySQL HeatWave to Amazon Kinesis](#)
- [Send data from MySQL HeatWave to Azure Event Hubs](#)
- [Replicate data from MySQL HeatWave to Google Cloud Storage](#)
- [Replicate Data from PostgreSQL to Google BigQuery](#)

MySQL quickstarts

- [Replicate data from OCI MySQL Database to Autonomous Data Warehouse](#)
- [Send Data from MySQL HeatWave to Azure Event Hubs](#)
- [Replicate data from MySQL HeatWave to Amazon Kinesis](#)
- [Replicate Data from MySQL HeatWave to Google Cloud Storage](#)

PostgreSQL quickstarts

- [Replicate Data from PostgreSQL to Autonomous Transaction Processing](#)
- [Replicate data from PostgreSQL to MySQL](#)
- [Replicate data from PostgreSQL to Snowflake](#)
- [Replicate data from PostgreSQL to Google BigQuery](#)

SQL Server quickstarts

- [Replicate data from Azure SQL Managed Instance to Autonomous Transaction Processing](#)

Stream Analytics quickstarts

- [Build a simple Stream Analytics pipeline](#)

- [Replicate data to Stream analytics](#)


Connect to OCI GoldenGate using a public IP

When you create an OCI GoldenGate deployment, you can select whether it's accessible through a public or private endpoint. This quickstart walks you through the steps to create a deployment with a public endpoint.

Overview

Note:

For Oracle Cloud Infrastructure Federal Government Cloud customers in regions that don't support DNS Zone Management, to enable a public endpoint, Create a Sev 1 Service Request titled, "Create DNS Entry for Public OCI GG Deployment," and include the correct deployment OCID. The support team will ensure that the appropriate DNS record gets created for your deployment.



Before you begin

To successfully complete this quickstart, you need:

- An Oracle Cloud Infrastructure account
- Access to OCI GoldenGate

Task 1: Create the deployment

1. In the Console navigation menu, click **Oracle Database**, and then select **GoldenGate**.
2. On the Deployments page, click **Create deployment**.
3. In the Create deployment panel, enter a name and optionally, a description.
4. From the Compartment dropdown, select a compartment in which to create the deployment.
5. Select one of the following options:
 - **Production:** Sets up a deployment with recommended defaults for a production environment. The minimum number of OCPUs is 4, with auto-scaling enabled.
 - **Development or testing:** Sets up a deployment with recommended defaults for a development or testing environment. The minimum number of OCPUs is 1.

6. For **OCPU count** enter the number of Oracle Compute units (OCPUs) to use.

 **Note:**

One OCPU is equivalent to 16gb of memory. For more information, see OCPU management and billing.

7. (Optional) Select **Auto scaling**.

 **Note:**

Auto scaling enables OCI GoldenGate to scale up to three times the number of OCPUs you specify for OCPU Count, up to 24 OCPUs. For example, if you specify your OCPU Count as 2 and enable Auto Scaling, then your deployment can scale up to 6 OCPUs. If you specify your OCPU Count as 20 and enable Auto Scaling, OCI GoldenGate can only scale up to 24 OCPUs.

8. From the **Subnet in <Compartment>** dropdown, select the subnet to which a private endpoint is created from the OCI GoldenGate service tenancy. This ensures that the deployment is always available over this subnet, as long as the policies for this subnet allow access. The private endpoint is only used to access the deployment console, and doesn't provide access to other resources in the subnet.

To select a subnet in a different compartment, click **Change compartment**.

 **Note:**

You can only select a private subnet when creating a deployment.

9. Select a license type.
10. (Optional) Click **Show advanced options** for network options and to add tags.
 - a. In the **Network** tab,
 - i. Select **Enable GoldenGate console public access** to include a public endpoint in addition to a private endpoint, and allow public access to the deployment console for users. If selected, OCI GoldenGate creates a load balancer in your tenancy to create a public IP. Select a subnet in the same VCN as this deployment in which to create the load balancer.

 **Note:**

The load balancer is a resource that comes with an additional cost. You can manage this resource, but ensure that you don't delete the load balancer while your deployment is still in use. [Learn more about load balancer pricing](#).

- ii. Select **Customize endpoint** to provide a private fully qualified domain name (FQDN) prefix that you'll use to access the private service console URL. You can also optionally upload an SSL/TLS certificate (.pem) and its corresponding private key, however, password protected certificates are not supported. It is your

responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.
A self-signed certificate is generated for you, if you don't provide one.

 **Note:**

It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.

- b. In the **Maintenance** tab:
 - i. Select **Customize maintenance window** to define the start of the maintenance window to upgrade the deployment.
 - ii. (Optional) For **Major release auto-upgrade period in days**, enter the number of days, between 0 and 365.
 - iii. (Optional) For **Bundle release auto-upgrade period in days**, enter the number of days, between 0 and 180 days.
 - iv. (Optional) For **Security patch auto-upgrade period in days**, enter the number of days, between 0 and 14 days.
 - v. Select **Enable interim release auto-upgrade**, and, optionally, enter the number of days.

 **Note:**

Learn more about scheduling upgrades.

- c. In the **Tags** tab, add tags to help track the resources within your tenancy. Click **+ Additional tag** to add more tags. [Learn more about tagging](#).
11. Click **Next**.
 12. For Deployment type, select **Data replication**.
 13. From the **Select a technology** dropdown, select one of the following technology types:
 - Oracle Database
 - Big Data
 - MySQL
 - PostgreSQL
 - Microsoft SQL Server

See what's supported to learn which databases and technologies you can use as OCI GoldenGate sources and targets.
 14. For **Version**, the latest version is automatically selected. Click **Change version** to select a different version.

 **Note:**

Learn more about versions.

15. For **GoldenGate instance name**, enter the name that the deployment will assign to the GoldenGate deployment instance upon creation.
16. For Credential store, select one of the following:
 - **OCI Identity and Access Management (OCI IAM)**, to enable users to log in to the the deployment console using their Oracle Cloud account (single sign on) in IAM (Identity and Access Management) enabled tenancies.

 **Note:**

Once you select IAM, you won't be able to switch to GoldenGate when you edit the deployment settings at a later time.

- **GoldenGate**, for GoldenGate to manage users.
 - a. Enter the **Administrator username**
 - b. Select a password secret in your compartment or click **Change compartment** to select one in a different compartment. You can also create a new password secret.

To create a new password secret:

 - i. Click **Create password secret**.
 - ii. In the Create secret panel, enter a name for the secret, and optionally, a description.
 - iii. Select a compartment from the **Compartment** dropdown in which to save your secret.
 - iv. Select a vault in the current compartment, or click **Change compartment** to select a vault in a different compartment.
 - v. Select an **Encryption key**.

 **Note:**

Only AES keys, Software protected keys, and HSM keys are supported. RSA and ECDSA keys are not supported for GoldenGate password secret keys.

- vi. Enter a password 8 to 30 characters in length, containing at least 1 uppercase, 1 lowercase, 1 numeric and 1 special character. The special characters must not be '\$', '^' or '?'.
- vii. Confirm the password.
- viii. Click **Create**.

Note:

You can manage GoldenGate users in the deployment console.
Learn more.

17. Click Create.

The deployment takes a few minutes to create. Once it becomes active, it's ready for you to use.

Task 2: Launch the console

1. From the Deployments list, select your deployment to view its details.
2. On the Deployment's details page, click **Launch Console**.

The deployment console opens in a new browser tab. Use the credentials you specified in Task 1 to log in to the deployment.

Learn more

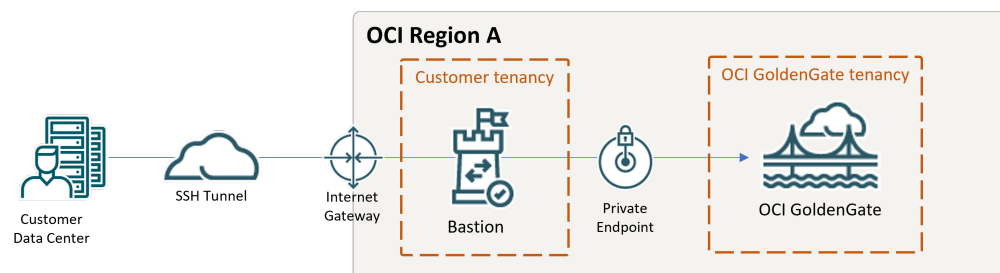
- Create deployments
- Manage deployments

Connect to Oracle Cloud Infrastructure GoldenGate using a private IP

Use OCI Bastion to secure access to your OCI GoldenGate Deployment Console.

Overview

OCI GoldenGate is only accessible using a private endpoint from within the OCI network, or through a bastion host that secures access to OCI resources. While this quickstart example uses OCI Bastion, it is possible for you to [use your own bastion](#). This quickstart includes both options, so you can choose the one that works best for you.



Before you begin

You must have the following in order to proceed:

- A free trial or paid Oracle Cloud Infrastructure account
- Access to OCI GoldenGate

- An OCI GoldenGate deployment in a private subnet and without a public endpoint
- For OCI Bastion:
 - Access to the service
 - Access to OCI Bastion or your own bastion on OCI Compute
- For your own bastion on OCI Compute:
 - Access to OCI Compute
 - Public and private subnets configured in each availability domain

 **Note:**

Oracle recommends creating a separate public subnet solely for bastion hosts to ensure that the appropriate security list is assigned to the correct host.

Option A: Use OCI Bastion

You can use OCI Bastion or use your own. This example uses OCI Bastion.

 **Note:**

For US Government Cloud with FedRAMP Authorization, you must use Option B. The OCI Bastion service is not currently available in these regions.

1. [Create a bastion](#). Ensure that you:
 - a. Use the same VCN as the target OCI GoldenGate deployment and subnet.

 **Note:**

The subnet can be the same as the OCI GoldenGate deployment or one that has access to the OCI GoldenGate subnet.

- b. Include the IP addresses of the machines used to connect to OCI Bastion in the **CIDR Block Allowlist**.
2. [Create a SSH port forwarding session](#).
 - a. For **IP Address**, enter the OCI GoldenGate deployment's private IP. You can find the private IP on the deployment's Details page.
 - b. For **Port**, enter 443.
 - c. Under **Add SSH Key**, provide the public key file of the SSH key pair to use for the session.
3. After the session is created, from the session's **Actions** (three dots) menu, select **Copy SSH Command**.

4. Paste the command into a text editor, and then replace the `<privateKey>` and `<localPort>` placeholders with the path to the private key and port 443.
5. Run the command using the command line interface to create the tunnel.
6. Open a web browser and go to `https://localhost`.

**Note:**

Ensure that you add an Ingress rule for the Bastion host in Private Subnet's security list. [Learn more](#).

Option B: Use your own bastion on OCI Compute

1. Create a compute instance in the public subnet of same VCN as the OCI GoldenGate deployment.

**Note:**

In this example, the public subnet CIDR is `10.0.0.0/24`. The same CIDR value will be used when you add an ingress rule to the private subnet security list.

2. Check the default security list for the public subnet:
 - a. From the Oracle Cloud console navigation menu, select **Networking**, and then **Virtual Cloud Networks**.
 - b. From the list of Virtual Cloud Networks, select your VCN to view its details.
 - c. Select the public subnet, and then select the security list to view its details. This security list must include a rule for SSH Access:

Stateless	Source	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows
No	0.0.0.0/0	TCP	All	22	N/A	TCP traffic for ports: 22 SSH Remote Login Protocol

If the security list doesn't include this rule, click **Add Ingress Rules** and complete the form using the values above.

3. Add an Ingress rule to the Private subnet security list to allow connectivity to OCI GoldenGate from the public subnet.
 - a. On the VCN details page, under **Subnets**, select the Private Subnet to view its details.
 - b. On the Private Subnet details page, under **Security Lists**, select the security list to view its details.
 - c. Under Ingress Rules, click **Add Ingress Rules**.
 - d. In the Add Ingress Rules dialog, complete the fields as follows, and then click **Add Ingress Rules**:

- i. For **Source Type**, select **CIDR**.
 - ii. For **Source CIDR**, enter the public subnet CIDR value (10.0.0.0/24).
 - iii. For **IP Protocol**, select **TCP**.
 - iv. For **Source Port Range**, enter 443.
4. (Windows users) Create a session to connect to the bastion host using PuTTY:
 - a. In the PuTTY Session configuration screen, enter the Compute instance's public IP for **Host Name**. You can leave 22 in as the value for **Port**.
 - b. Under the Connection category, expand **SSH**, click **Auth**, and then click Browse to locate the private you used to create the Compute instance.
 - c. Click **Tunnels** in the Category panel, enter 443 for Source port, and <deployment-hostname>:443 for Destination.
 - d. (Optional) Return to the Session category and Save the session details.
 - e. Click **Open** to connect.
5. (Linux users) Create a session to connect to the bastion host using the command line:



```
ssh -i <private-ssh-key> opc@<compute-public-ip> -L 443:<deployment-hostname>:443 -N
```
6. After successfully connected, open a browser window and enter `https://localhost` in the address bar. You're brought to the OCI GoldenGate deployment console.

Connect to OCI GoldenGate using a public load balancer

Learn to create and configure a public load balancer in your tenancy to access a private OCI GoldenGate deployment.

Overview

When you create an OCI GoldenGate deployment, you can enable or disable the deployment's public endpoint. Because the OCI GoldenGate Public Endpoint is managed by the OCI GoldenGate service tenancy, it's not possible for you to create network security group (NSG) rules from your customer tenancy.



Before you begin

You must have the following in order to proceed:

- A free or paid Oracle Cloud Infrastructure account
- Access to OCI GoldenGate and Networking services

- Access to DNS service or third-party DNS management system, such as GoDaddy

Task 1: Create a certificate bundle

Create a certificate bundle that includes the public certificate, the corresponding private key, and any associated Certificate Authority (CA) certificates. For more information, see [SSL Certificate for Load Balancers](#).

Task 2: Create a deployment

1. In the Console navigation menu, click **Oracle Database**, and then select **GoldenGate**.
2. On the Deployments page, click **Create deployment**.
3. In the Create deployment panel, enter a name and optionally, a description.
4. From the Compartment dropdown, select a compartment in which to create the deployment.
5. Select one of the following options:
 - **Production**: Sets up a deployment with recommended defaults for a production environment. The minimum number of OCPUs is 4, with auto-scaling enabled.
 - **Development or testing**: Sets up a deployment with recommended defaults for a development or testing environment. The minimum number of OCPUs is 1.
6. For **OCPU count** enter the number of Oracle Compute units (OCPUs) to use.

 **Note:**

One OCPU is equivalent to 16gb of memory. For more information, see OCPU management and billing.

7. (Optional) Select **Auto scaling**.

 **Note:**

Auto scaling enables OCI GoldenGate to scale up to three times the number of OCPUs you specify for OCPU Count, up to 24 OCPUs. For example, if you specify your OCPU Count as 2 and enable Auto Scaling, then your deployment can scale up to 6 OCPUs. If you specify your OCPU Count as 20 and enable Auto Scaling, OCI GoldenGate can only scale up to 24 OCPUs.

8. From the **Subnet in <Compartment>** dropdown, select the subnet to which a private endpoint is created from the OCI GoldenGate service tenancy. This ensures that the deployment is always available over this subnet, as long as the policies for this subnet allow access. The private endpoint is only used to access the deployment console, and doesn't provide access to other resources in the subnet.

To select a subnet in a different compartment, click **Change compartment**.

 **Note:**

You can only select a private subnet when creating a deployment.

9. Select a license type.
10. (Optional) Click **Show advanced options** for network options and to add tags.
 - a. In the **Network** tab,
 - i. Select **Enable GoldenGate console public access** to include a public endpoint in addition to a private endpoint, and allow public access to the deployment console for users. If selected, OCI GoldenGate creates a load balancer in your tenancy to create a public IP. Select a subnet in the same VCN as this deployment in which to create the load balancer.

 **Note:**

The load balancer is a resource that comes with an additional cost. You can manage this resource, but ensure that you don't delete the load balancer while your deployment is still in use. [Learn more about load balancer pricing](#).

- ii. Select **Customize endpoint** to provide a private fully qualified domain name (FQDN) prefix that you'll use to access the private service console URL. You can also optionally upload an SSL/TLS certificate (.pem) and its corresponding private key, however, password protected certificates are not supported. It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.
A self-signed certificate is generated for you, if you don't provide one.

 **Note:**

It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.

- b. In the **Maintenance** tab:
 - i. Select **Customize maintenance window** to define the start of the maintenance window to upgrade the deployment.
 - ii. (Optional) For **Major release auto-upgrade period in days**, enter the number of days, between 0 and 365.
 - iii. (Optional) For **Bundle release auto-upgrade period in days**, enter the number of days, between 0 and 180 days.
 - iv. (Optional) For **Security patch auto-upgrade period in days**, enter the number of days, between 0 and 14 days.
 - v. Select **Enable interim release auto-upgrade**, and, optionally, enter the number of days.

 **Note:**

Learn more about scheduling upgrades.

- c. In the **Tags** tab, add tags to help track the resources within your tenancy. Click **+ Additional tag** to add more tags. [Learn more about tagging.](#)
11. Click **Next**.
12. For Deployment type, select **Data replication**.
13. From the **Select a technology** dropdown, select one of the following technology types:
 - Oracle Database
 - Big Data
 - MySQL
 - PostgreSQL
 - Microsoft SQL Server

See what's supported to learn which databases and technologies you can use as OCI GoldenGate sources and targets.

14. For **Version**, the latest version is automatically selected. Click **Change version** to select a different version.

 **Note:**

Learn more about versions.

15. For **GoldenGate instance name**, enter the name that the deployment will assign to the GoldenGate deployment instance upon creation.
16. For Credential store, select one of the following:
 - **OCI Identity and Access Management (OCI IAM)**, to enable users to log in to the the deployment console using their Oracle Cloud account (single sign on) in IAM (Identity and Access Management) enabled tenancies.

 **Note:**

Once you select IAM, you won't be able to switch to GoldenGate when you edit the deployment settings at a later time.

- **GoldenGate**, for GoldenGate to manage users.
 - a. Enter the **Administrator username**
 - b. Select a password secret in your compartment or click **Change compartment** to select one in a different compartment. You can also create a new password secret.

To create a new password secret:

 - i. Click **Create password secret**.

- ii. In the Create secret panel, enter a name for the secret, and optionally, a description.
- iii. Select a compartment from the **Compartment** dropdown in which to save your secret.
- iv. Select a vault in the current compartment, or click **Change compartment** to select a vault in a different compartment.
- v. Select an **Encryption key**.

 **Note:**

Only AES keys, Software protected keys, and HSM keys are supported. RSA and ECDSA keys are not supported for GoldenGate password secret keys.

- vi. Enter a password 8 to 30 characters in length, containing at least 1 uppercase, 1 lowercase, 1 numeric and 1 special character. The special characters must not be '\$', '^' or '?'.
- vii. Confirm the password.
- viii. Click **Create**.

 **Note:**

You can manage GoldenGate users in the deployment console. Learn more.

17. Click **Create**.

Task 3: Create the load balancer

To create a load balancer with SSL:

1. In the OCI Console navigation menu, select **Networking**, and then click **Load Balancers**.
2. On the **Load Balancers** page, click **Create Load Balancer**.
3. In the **Select Load Balancer Type** dialog, select **Load Balancer**, and then click **Create Load Balancer**.
4. On the **Add Details** page, complete the following fields, and then click **Next**:
 - a. For **Load Balancer Name**, enter a name.
 - b. For **Visibility**, select either **Public** or **Private**.
 - c. For **Assign a public IP address**, select **Reserved IP**.
 - d. For **Shapes**, select **Dynamic** and then move the selector from Small to Micro.
 - e. For **Choose Networking**, select your VCN and subnet from their respective dropdowns.
5. On the **Choose Backends** page, complete the following fields, and then click **Next**:
 - a. For **Specify a Load Balancing Policy**, select **Weighted Round Robin**.

- b. Under **Specify Health Check Policy**, select **TCP** from the Protocol dropdown, and then enter **443** for **Port**.
 - c. Leave **SSL** unchecked.
6. On the **Configure Listener** page, completed the following fields, and then click **Next**:
 - a. For **Specify the type of traffic your listener handles**, select **HTTPS**.
 - b. For **Specify the port your listener monitors for ingress traffic**, ensure that **443** is displayed.
 - c. For **SSL Certificate**, drag-and-drop or select the SSL Certificate (.cer).
 - d. Select **Specify CA Certificate** and then drag-and-drop or select the CA Certificate (.crt).
 - e. Select **Specify Private Key**, and then drag-and-drop or select the Private Key File.
7. On the **Managing Logging** page, complete the following fields, disable **Error Logs**, and then click **Submit**.
8. On the **Load Balancer Details** page, under **Resources**, click **Backend Sets**.
9. Under **Backend Sets**, select the backend set displayed in the list, and then click **Edit**.
10. In the **Edit Backend Set** panel, select **Use SSL**, ensure that your certificate is selected, and then click **Save Changes**.
11. On the **Backend Sets Details** page, under **Resources**, click **Backends**, and then click **Add Backends**.
12. In the **Add Backends** panel, select **IP Addresses**, enter the OCI GoldenGate deployment's Private IP Address (from Step 2) for **IP Address**, and **443** in for **Port**, and then click **Add**.
13. In the breadcrumb, click **Load Balancer Details**, and then copy the **IP Address**.

You can use a web browser to access this IP address, verify the certificate is the digitally signed certificate that you uploaded, and access the OCI GoldenGate Deployment Console. Next, you'll create a DNS record for the Load Balancer's IP.

Task 4: Create a DNS record

Create a DNS record for the Load Balancer's Public IP in a DNS management system.

You can use [Oracle Cloud Infrastructure DNS Management](#) or any public DNS management system.

After a few minutes, verify that you can access the OCI GoldenGate Deployment Console through the domain you created.

Task 5: Create OCI Network Security Rules to allow/deny ingress

1. From the OCI Console navigation menu (hamburger icon), click **Networking**, then **Virtual Cloud Networks**.
2. From the **Virtual Cloud Networks** list, select your VCN.

3. On the **VCN Details** page, select your subnet.
4. On the **Subnet Details** page, copy the **IPv4 CIDR Block** value, and then click **Default Security List for <VCN>** under **Security Lists**.
5. On the **Default Security Lists Details** page, under **Ingress Rules**, locate the ingress rule for TCP that is currently open for all source and destination port ranges, and then select **Edit** from its **Actions** (ellipsis) menu.
6. In the **Edit Ingress Rule** dialog, replace the **Source CIDR** value with the IPv4 CIDR Block value copied from Step 5d, and then click **Save Changes**.

Wait a few minutes for the changes to take effect.

7. Click **Add Ingress Rule**, and then replace the **Source CIDR** value with an IP address range that includes the Load Balancer's IP address, and then click **Add Ingress Rules**.

You can also add an ingress rule for the IP address of your local machine to verify that the routing rules are in effect.

Learn more

- [Oracle Cloud Infrastructure Load Balancing](#)

Oracle Database quickstarts

Common use cases using Oracle Databases as OCI GoldenGate sources or targets.

Articles in this section:

- [Replicate data between cloud databases in the same region](#)
- [Replicate data between cloud databases in different regions](#)
- [Replicate data between cloud databases in different region with VCN peering](#)
- [Configure bidirectional replication](#)
- [Send data from Oracle GoldenGate to OCI GoldenGate](#)
- [Send data from OCI GoldenGate to Oracle GoldenGate](#)

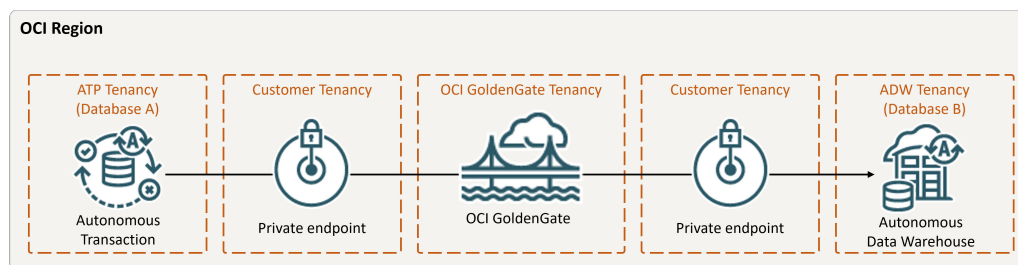
Replicate data between cloud databases in the same region

Learn how to set up Oracle Cloud Infrastructure GoldenGate to replicate data between two Autonomous Databases.

Overview

Oracle Cloud Infrastructure GoldenGate lets you to replicate supported databases within the same region. The following steps guide you through how to instantiate a target database using Oracle Data Pump and replicate data from the source to the target.

This quickstart is also available as LiveLab: [View the workshop](#).



Before you begin

You must have the following in order to proceed:

- An existing source database
- An existing target database
- The source and target database must be in a single tenancy, in the same region
- If you need sample data, download [Archive.zip](#), and then follow the instructions in [Lab 1, Task 3: Load the ATP schema](#).

Task 1: Set up the environment

1. Create a deployment.
2. Create connections.
3. Create an association between connections and deployments.
4. Run the following query to ensure that `support_mode=FULL` for all tables in the source database:

```
select * from DBA_GOLDENGATE_SUPPORT_MODE where owner =
'SRC_OCIGLL';
```

5. Enable supplemental logging:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA
```

Task 2: Create the Integrated Extract

An Integrated Extract captures ongoing changes to source database.

1. On the deployment Details page, click **Launch console**.
2. Add an Extract.

Note:

See additional extract parameter options for more information about parameters that you can use to specify source tables.

- On the **Extract Parameters** page, append the following lines under `EXTTRAIL` `<extract-name>`:

```
-- Capture DDL operations for listed schema tables
ddl include mapped

-- Add step-by-step history of
-- to the report file. Very useful when troubleshooting.
ddloptions report

-- Write capture stats per table to the report file daily.
report at 00:01

-- Rollover the report file weekly. Useful when IE runs
-- without being stopped/started for long periods of time to
-- keep the report files from becoming too large.
reportrollover at 00:01 on Sunday

-- Report total operations captured, and operations per second
-- every 10 minutes.
reportcount every 10 minutes, rate

-- Table list for capture
table SRC_OCIGLL.*;
```

3. Check for long running transactions:

- Run the following script on your source database:

```
select start_scn, start_time from gv$transaction where start_scn <
(select max(start_scn) from dba_capture);
```

If the query returns any rows, then you must locate the transaction's SCN and then either commit or rollback the transaction.

Task 3: Export data using Oracle Data Pump (ExpDP)

Use Oracle Data Pump (ExpDP) to export data from the source database to Oracle Object Store.

1. Create an Oracle Object Store bucket.

Take note of the namespace and bucket name for use with the Export and Import scripts.

2. Create an Auth Token, and then copy and paste the token string to a text editor for later use.

3. Create a credential in your source database, replacing the `<user-name>` and `<token>` with your Oracle Cloud account username and the token string you created in the previous step:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'ADB_OBJECTSTORE',
    username => '<user-name>',
```

```

        password => '<token>'
    );
END;
```

4. Run the following script in your source database to create the Export Data job. Ensure that you replace the <region>, <namespace>, and <bucket-name> in Object Store URI accordingly. SRC_OCIGLL.dmp is a file that will be created when this script runs.

```

DECLARE
ind NUMBER;           -- Loop index
h1 NUMBER;           -- Data Pump job handle
percent_done NUMBER; -- Percentage of job complete
job_state VARCHAR2(30); -- To keep track of job state
le ku$_LogEntry;     -- For WIP and error messages
js ku$_JobStatus;    -- The job status from get_status
jd ku$_JobDesc;      -- The job description from get_status
sts ku$_Status;      -- The status object returned by get_status

BEGIN
-- Create a (user-named) Data Pump job to do a schema export.
h1 :=
DBMS_DATAPUMP.OPEN('EXPORT','SCHEMA',NULL,'SRC_OCIGLL_EXPORT','LATEST');

-- Specify a single dump file for the job (using the handle just
returned
-- and a directory object, which must already be defined and
accessible
-- to the user running this procedure.
DBMS_DATAPUMP.ADD_FILE(h1,'https://
objectstorage.<region>.oraclecloud.com/n/<namespace>/b/<bucket-
name>/o/
SRC_OCIGLL.dmp','ADB_OBJECTSTORE','100MB',DBMS_DATAPUMP.KU$_FILE_TY
PE_URIDUMP_FILE,1);

-- A metadata filter is used to specify the schema that will be
exported.
DBMS_DATAPUMP.METADATA_FILTER(h1,'SCHEMA_EXPR','IN
(''SRC_OCIGLL'')');

-- Start the job. An exception will be generated if something is
not set up properly.
DBMS_DATAPUMP.START_JOB(h1);

-- The export job should now be running. In the following loop, the
job
-- is monitored until it completes. In the meantime, progress
information is displayed.
percent_done := 0;
job_state := 'UNDEFINED';
while (job_state != 'COMPLETED') and (job_state != 'STOPPED') loop
    dbms_datapump.get_status(h1,dbms_datapump.ku$_status_job_error +
dbms_datapump.ku$_status_job_status +
dbms_datapump.ku$_status_wip,-1,job_state,sts);
```

```

        js := sts.job_status;

-- If the percentage done changed, display the new value.
if js.percent_done != percent_done
then
    dbms_output.put_line('*** Job percent done = ' ||
to_char(js.percent_done));
    percent_done := js.percent_done;
end if;

-- If any work-in-progress (WIP) or error messages were received for the
job, display them.
if (bitand(sts.mask,dbms_datapump.ku$_status_wip) != 0)
then
    le := sts.wip;
else
    if (bitand(sts.mask,dbms_datapump.ku$_status_job_error) != 0)
    then
        le := sts.error;
    else
        le := null;
    end if;
end if;
if le is not null
then
    ind := le.FIRST;
    while ind is not null loop
        dbms_output.put_line(le(ind).LogText);
        ind := le.NEXT(ind);
    end loop;
end if;
end loop;

-- Indicate that the job finished and detach from it.
dbms_output.put_line('Job has completed');
dbms_output.put_line('Final job state = ' || job_state);
dbms_datapump.detach(h1);
END;
```

Task 4: Instantiate the target database using Oracle Data Pump (ImpDP)

Use Oracle Data Pump (ImpDP) to import data into the target database from the SRC_OCIGLLL.dmp that was exported from the source database.

1. Create a credential in your target database to access Oracle Object Store (using the same information in the preceding section).

```

BEGIN
    DBMS_CLOUD.CREATE_CREDENTIAL(
        credential_name => 'ADB_OBJECTSTORE',
        username => '<user-name>',
        password => '<token>'
```

```
);  
END;
```

2. Run the following script in your target database to import data from the SRC_OCIGLL.dmp. Ensure that you replace the <region>, <namespace>, and <bucket-name> in Object Store URI accordingly:

```
DECLARE  
ind NUMBER; -- Loop index  
h1 NUMBER; -- Data Pump job handle  
percent_done NUMBER; -- Percentage of job complete  
job_state VARCHAR2(30); -- To keep track of job state  
le ku$_LogEntry; -- For WIP and error messages  
js ku$_JobStatus; -- The job status from get_status  
jd ku$_JobDesc; -- The job description from get_status  
sts ku$_Status; -- The status object returned by get_status  
BEGIN  
  
-- Create a (user-named) Data Pump job to do a "full" import  
(everything  
-- in the dump file without filtering).  
h1 :=  
DBMS_DATAPUMP.OPEN('IMPORT','FULL',NULL,'SRCMIRROR_OCIGLL_IMPORT');  
  
-- Specify the single dump file for the job (using the handle just  
returned)  
-- and directory object, which must already be defined and  
accessible  
-- to the user running this procedure. This is the dump file  
created by  
-- the export operation in the first example.  
  
DBMS_DATAPUMP.ADD_FILE(h1,'https://  
objectstorage.<region>.oraclecloud.com/n/<namespace>/b/<bucket-  
name>/o/  
SRC_OCIGLL.dmp','ADB_OBJECTSTORE',null,DBMS_DATAPUMP.KU$_FILE_TYPE_  
URIDUMP_FILE);  
  
-- A metadata remap will map all schema objects from SRC_OCIGLL to  
SRCMIRROR_OCIGLL.  
DBMS_DATAPUMP.METADATA_REMAP(h1,'REMAP_SCHEMA','SRC_OCIGLL','SRCMIR  
ROR_OCIGLL');  
  
-- If a table already exists in the destination schema, skip it  
(leave  
-- the preexisting table alone). This is the default, but it does  
not hurt  
-- to specify it explicitly.  
DBMS_DATAPUMP.SET_PARAMETER(h1,'TABLE_EXISTS_ACTION','SKIP');  
  
-- Start the job. An exception is returned if something is not set  
up properly.  
DBMS_DATAPUMP.START_JOB(h1);
```

```
-- The import job should now be running. In the following loop, the job is
-- monitored until it completes. In the meantime, progress information is
-- displayed. Note: this is identical to the export example.
percent_done := 0;
job_state := 'UNDEFINED';
while (job_state != 'COMPLETED') and (job_state != 'STOPPED') loop
    dbms_datapump.get_status(h1,
        dbms_datapump.ku$_status_job_error +
        dbms_datapump.ku$_status_job_status +
        dbms_datapump.ku$_status_wip,-1,job_state,sts);
    js := sts.job_status;

    -- If the percentage done changed, display the new value.
    if js.percent_done != percent_done
    then
        dbms_output.put_line('*** Job percent done = ' ||
            to_char(js.percent_done));
        percent_done := js.percent_done;
    end if;

    -- If any work-in-progress (WIP) or Error messages were received for
    the job, display them.
    if (bitand(sts.mask,dbms_datapump.ku$_status_wip) != 0)
    then
        le := sts.wip;
    else
        if (bitand(sts.mask,dbms_datapump.ku$_status_job_error) != 0)
        then
            le := sts.error;
        else
            le := null;
        end if;
    end if;
    if le is not null
    then
        ind := le.FIRST;
        while ind is not null loop
            dbms_output.put_line(le(ind).LogText);
            ind := le.NEXT(ind);
        end loop;
    end if;
end loop;

-- Indicate that the job finished and gracefully detach from it.
dbms_output.put_line('Job has completed');
dbms_output.put_line('Final job state = ' || job_state);
dbms_datapump.detach(h1);
END;
```

Task 5: Add and run a Non-integrated Replicat

1. Add and run a Replicat.

- On the **Parameter File** screen, replace `MAP *.*`, `TARGET *.*`; with the following script:

```
-- Capture DDL operations for listed schema tables
--
ddl include mapped
--
-- Add step-by-step history of ddl operations captured
-- to the report file. Very useful when troubleshooting.
--
ddloptions report
--
-- Write capture stats per table to the report file daily.
--
report at 00:01
--
-- Rollover the report file weekly. Useful when PR runs
-- without being stopped/started for long periods of time to
-- keep the report files from becoming too large.
--
reportrollover at 00:01 on Sunday
--
-- Report total operations captured, and operations per second
-- every 10 minutes.
--
reportcount every 10 minutes, rate
--
-- Table map list for apply
--
DBOPTIONS ENABLE_INSTANTIATION_FILTERING;
MAP SRC_OCIGLL.*, TARGET SRCMIRROR_OCIGLL.*;
```

 **Note:**

`DBOPTIONS ENABLE_INSTANTIATION_FILTERING` enables CSN filtering on tables imported using Oracle Data Pump. For more information, see [DBOPTIONS Reference](#).

2. Perform Inserts to the source database:
 - a. Return to the Oracle Cloud console and use the navigation menu to navigate back to **Oracle Database, Autonomous Transaction Processing**, and then **SourceATP**.
 - b. On the Source ATP Details page, click **Tools**, and then **Database actions**.
 - c. Use the Source ATP database credentials in the Workshop details to log in to Database actions, and then click **SQL**.
 - d. Enter the following inserts, and then click **Run Script**:

```
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1000,'Houston',20,743113);
Insert into SRC_OCIGLL.SRC_CITY
```

```
(CITY_ID,CITY,REGION_ID,POPULATION) values (1001,'Dallas',20,822416);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1002,'San Francisco',21,157574);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1003,'Los Angeles',21,743878);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1004,'San Diego',21,840689);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1005,'Chicago',23,616472);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1006,'Memphis',23,580075);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1007,'New York City',22,124434);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1008,'Boston',22,275581);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1009,'Washington D.C.',22,688002);
```

- e. In the OCI GoldenGate Deployment Console, click the **Extract name (UAEXT)**, and then click **Statistics**. Verify that **SRC_OCIGLL.SRC_CITY** is listed with 10 inserts.
- f. Go back to the Overview screen, click the **Replicat name (REP)**, and then click **Statistics**. Verify that **SRCMIRROR_OCIGLL.SRC_CITY** is listed with 10 inserts

Task 6: Monitor and maintain processes

1. Monitor the replication process.
2. Manage Trail files.

Configure bidirectional replication between two cloud databases in the same region

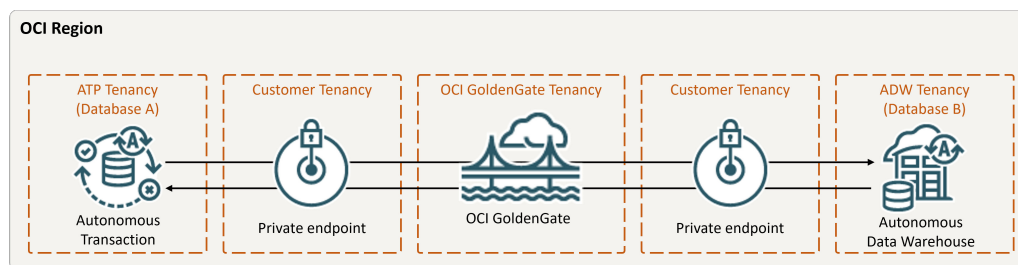
After you set up unidirectional replication, there's just a few extra steps to replicate data in the opposite direction. This quickstart example uses Autonomous Transaction Processing and Autonomous Data Warehouse as its two cloud databases.

Before you begin

You must have two existing databases in the same tenancy and region in order to proceed with this quickstart. If you need sample data, download [Archive.zip](#), and then follow the instructions in [Lab 1, Task 3: Load the ATP schema](#)

Overview

The following steps guide you through how to instantiate a target database using Oracle Data Pump and set up bidirectional (two-way) replication between two databases in the same region.



Task 1: Set up the environment

1. Create a deployment.
2. Create connections to your databases.
3. Assign the connections to the deployment.
4. Enable supplemental logging:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA
```

5. Run the following query to ensure that `support_mode=FULL` for all tables in the source database:

```
select * from DBA_GOLDENGATE_SUPPORT_MODE where owner =
'SRC_OCIGLL';
```

6. Run the following query on Database B to ensure that `support_mode=FULL` for all tables in the database:

```
select * from DBA_GOLDENGATE_SUPPORT_MODE where owner =
'SRCMIRROR_OCIGLL';
```

Task 2: Add transaction information and a checkpoint table for both databases

In the OCI GoldenGate deployment console, go to the Configuration screen for the Administration Service, and then complete the following:

1. Add transaction information on Database A and B:
 - a. For Database A, enter `SRC_OCIGLL` for Schema Name.
 - b. For Database B, enter `SRCMIRROR_OCIGLL` for Schema Name.

Note:

The schema names should be unique, and match your database schema names if you're using a different dataset from this example.

2. Create a Checkpoint table for Database A and B:
 - a. For Database A, enter `"SRC_OCIGLL"."ATP_CHECKTABLE"` for Checkpoint Table.

- b. For Database B, enter "SRCMIRROR_OCIGLL"."CHECKTABLE" for Checkpoint Table.

Task 3: Create the Integrated Extract

An Integrated Extract captures ongoing changes to source database.

1. On the deployment Details page, click **Launch console**.
2. Add and run an Integrated Extract.

 **Note:**

See additional extract parameter options for more information about parameters that you can use to specify source tables.

- On the Extract Parameters page, append the following lines under `EXTTRAIL` `<extract-name>`:

```
-- Capture DDL operations for listed schema tables
ddl include mapped

-- Add step-by-step history of
-- to the report file. Very useful when troubleshooting.
ddloptions report

-- Write capture stats per table to the report file daily.
report at 00:01

-- Rollover the report file weekly. Useful when IE runs
-- without being stopped/started for long periods of time to
-- keep the report files from becoming too large.
reportrollover at 00:01 on Sunday

-- Report total operations captured, and operations per second
-- every 10 minutes.
reportcount every 10 minutes, rate

-- Table list for capture
table SRC_OCIGLL.*;

-- Exclude changes made by GGADMIN
tranlogoptions excludeuser ggadmin
```

 **Note:**

`tranlogoptions excludeuser ggadmin` avoids recapturing transactions applied by `ggadmin` in bidirectional replication scenarios.

3. Check for long running transactions:

- Run the following script on your source database:

```
select start_scn, start_time from gv$transaction where start_scn
< (select max(start_scn) from dba_capture);
```

If the query returns any rows, then you must locate the transaction's SCN and then either commit or rollback the transaction.

Task 4: Export data using Oracle Data Pump (ExpDP)

Use Oracle Data Pump (ExpDP) to export data from the source database to Oracle Object Store.

1. [Create an Oracle Object Store bucket.](#)

Take note of the namespace and bucket name for use with the Export and Import scripts.

2. [Create an Auth Token](#), and then copy and paste the token string to a text editor for later use.
3. Create a credential in your source database, replacing the `<user-name>` and `<token>` with your Oracle Cloud account username and the token string you created in the previous step:

```
BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'ADB_OBJECTSTORE',
    username => '<user-name>',
    password => '<token>'
  );
END;
```

4. Run the following script in your source database to create the Export Data job. Ensure that you replace the `<region>`, `<namespace>`, and `<bucket-name>` in Object Store URI accordingly. `SRC_OCIGLL.dmp` is a file that will be created when this script runs.

```
DECLARE
  ind NUMBER;           -- Loop index
  h1 NUMBER;           -- Data Pump job handle
  percent_done NUMBER; -- Percentage of job complete
  job_state VARCHAR2(30); -- To keep track of job state
  le ku$_LogEntry;     -- For WIP and error messages
  js ku$_JobStatus;    -- The job status from get_status
  jd ku$_JobDesc;      -- The job description from get_status
  sts ku$_Status;      -- The status object returned by get_status

BEGIN
  -- Create a (user-named) Data Pump job to do a schema export.
  h1 :=
  DBMS_DATAPUMP.OPEN('EXPORT', 'SCHEMA', NULL, 'SRC_OCIGLL_EXPORT', 'LATEST');
END;
```

```
-- Specify a single dump file for the job (using the handle just returned
-- and a directory object, which must already be defined and accessible
-- to the user running this procedure.
DBMS_DATAPUMP.ADD_FILE(h1,'https://
objectstorage.<region>.oraclecloud.com/n/<namespace>/b/<bucket-name>/o/
SRC_OCIGLL.dmp','ADB_OBJECTSTORE','100MB',DBMS_DATAPUMP.KU$_FILE_TYPE_URI
DUMP_FILE,1);

-- A metadata filter is used to specify the schema that will be exported.
DBMS_DATAPUMP.METADATA_FILTER(h1,'SCHEMA_EXPR','IN (''SRC_OCIGLL'')');

-- Start the job. An exception will be generated if something is not set
up properly.
DBMS_DATAPUMP.START_JOB(h1);

-- The export job should now be running. In the following loop, the job
-- is monitored until it completes. In the meantime, progress information
is displayed.
percent_done := 0;
job_state := 'UNDEFINED';
while (job_state != 'COMPLETED') and (job_state != 'STOPPED') loop
    dbms_datapump.get_status(h1,dbms_datapump.ku$_status_job_error +
    dbms_datapump.ku$_status_job_status +
    dbms_datapump.ku$_status_wip,-1,job_state,sts);
    js := sts.job_status;

-- If the percentage done changed, display the new value.
if js.percent_done != percent_done
then
    dbms_output.put_line('*** Job percent done = ' ||
to_char(js.percent_done));
    percent_done := js.percent_done;
end if;

-- If any work-in-progress (WIP) or error messages were received for the
job, display them.
if (bitand(sts.mask,dbms_datapump.ku$_status_wip) != 0)
then
    le := sts.wip;
else
    if (bitand(sts.mask,dbms_datapump.ku$_status_job_error) != 0)
then
        le := sts.error;
    else
        le := null;
    end if;
end if;
if le is not null
then
    ind := le.FIRST;
    while ind is not null loop
        dbms_output.put_line(le(ind).LogText);
        ind := le.NEXT(ind);
    end loop;
end if;
```

```

end loop;

-- Indicate that the job finished and detach from it.
dbms_output.put_line('Job has completed');
dbms_output.put_line('Final job state = ' || job_state);
dbms_datapump.detach(h1);
END;

```

Task 5: Instantiate the target database using Oracle Data Pump (ImpDP)

Use Oracle Data Pump (ImpDP) to import data into the target database from the SRC_OCIGLL.dmp that was exported from the source database.

1. Create a credential in your target database to access Oracle Object Store (using the same information in the preceding section).

```

BEGIN
  DBMS_CLOUD.CREATE_CREDENTIAL(
    credential_name => 'ADB_OBJECTSTORE',
    username => '<user-name>',
    password => '<token>'
  );
END;

```

2. Run the following script in your target database to import data from the SRC_OCIGLL.dmp. Ensure that you replace the <region>, <namespace>, and <bucket-name> in Object Store URI accordingly:

```

DECLARE
ind NUMBER; -- Loop index
h1 NUMBER; -- Data Pump job handle
percent_done NUMBER; -- Percentage of job complete
job_state VARCHAR2(30); -- To keep track of job state
le ku$_LogEntry; -- For WIP and error messages
js ku$_JobStatus; -- The job status from get_status
jd ku$_JobDesc; -- The job description from get_status
sts ku$_Status; -- The status object returned by get_status
BEGIN

-- Create a (user-named) Data Pump job to do a "full" import
(everything
-- in the dump file without filtering).
h1 :=
DBMS_DATAPUMP.OPEN('IMPORT','FULL',NULL,'SRCMIRROR_OCIGLL_IMPORT');

-- Specify the single dump file for the job (using the handle just
returned)
-- and directory object, which must already be defined and
accessible
-- to the user running this procedure. This is the dump file
created by
-- the export operation in the first example.

DBMS_DATAPUMP.ADD_FILE(h1,'https://

```

```
objectstorage.<region>.oraclecloud.com/n/<namespace>/b/<bucket-name>/o/  
SRC_OCIGLL.dmp', 'ADB_OBJECTSTORE', null, DBMS_DATAPUMP.KU$_FILE_TYPE_URIDUM  
P_FILE);  
  
-- A metadata remap will map all schema objects from SRC_OCIGLL to  
SRCMIRROR_OCIGLL.  
DBMS_DATAPUMP.METADATA_REMAP(h1, 'REMAP_SCHEMA', 'SRC_OCIGLL', 'SRCMIRROR_OC  
IGLL');  
  
-- If a table already exists in the destination schema, skip it (leave  
-- the preexisting table alone). This is the default, but it does not hurt  
-- to specify it explicitly.  
DBMS_DATAPUMP.SET_PARAMETER(h1, 'TABLE_EXISTS_ACTION', 'SKIP');  
  
-- Start the job. An exception is returned if something is not set up  
properly.  
DBMS_DATAPUMP.START_JOB(h1);  
  
-- The import job should now be running. In the following loop, the job is  
-- monitored until it completes. In the meantime, progress information is  
-- displayed. Note: this is identical to the export example.  
percent_done := 0;  
job_state := 'UNDEFINED';  
while (job_state != 'COMPLETED') and (job_state != 'STOPPED') loop  
    dbms_datapump.get_status(h1,  
        dbms_datapump.ku$_status_job_error +  
        dbms_datapump.ku$_status_job_status +  
        dbms_datapump.ku$_status_wip, -1, job_state, sts);  
    js := sts.job_status;  
  
    -- If the percentage done changed, display the new value.  
    if js.percent_done != percent_done  
    then  
        dbms_output.put_line('*** Job percent done = ' ||  
            to_char(js.percent_done));  
        percent_done := js.percent_done;  
    end if;  
  
    -- If any work-in-progress (WIP) or Error messages were received for  
    the job, display them.  
    if (bitand(sts.mask, dbms_datapump.ku$_status_wip) != 0)  
    then  
        le := sts.wip;  
    else  
        if (bitand(sts.mask, dbms_datapump.ku$_status_job_error) != 0)  
        then  
            le := sts.error;  
        else  
            le := null;  
        end if;  
    end if;  
    if le is not null  
    then  
        ind := le.FIRST;
```

```
        while ind is not null loop
            dbms_output.put_line(le(ind).LogText);
            ind := le.NEXT(ind);
        end loop;
    end if;
end loop;

-- Indicate that the job finished and gracefully detach from it.
dbms_output.put_line('Job has completed');
dbms_output.put_line('Final job state = ' || job_state);
dbms_datapump.detach(h1);
END;
```

Task 6: Add and run a Non-integrated Replicat

1. Add and run a Replicat.

- On the **Parameter File** screen, replace `MAP *.*`, `TARGET *.*`; with the following script:

```
-- Capture DDL operations for listed schema tables
--
ddl include mapped
--
-- Add step-by-step history of ddl operations captured
-- to the report file. Very useful when troubleshooting.
--
ddloptions report
--
-- Write capture stats per table to the report file daily.
--
report at 00:01
--
-- Rollover the report file weekly. Useful when PR runs
-- without being stopped/started for long periods of time to
-- keep the report files from becoming too large.
--
reportrollover at 00:01 on Sunday
--
-- Report total operations captured, and operations per second
-- every 10 minutes.
--
reportcount every 10 minutes, rate
--
-- Table map list for apply
--
DBOPTIONS ENABLE_INSTANTIATION_FILTERING;
MAP SRC_OCIGLL.*, TARGET SRCMIRROR_OCIGLL.*;
```

 **Note:**

`DBOPTIONS ENABLE_INSTANTIATION_FILTERING` enables CSN filtering on tables imported using Oracle Data Pump. For more information, see [DBOPTIONS Reference](#).

2. Perform some changes on Database A to see them replicated to Database B.

Task 7: Configure replication from Database B to Database A

Tasks 1 through 6 established replication from Database A to Database B. The following steps sets up replication from Database B to Database A.

1. Add and run an Extract on Database B. On the extract parameters page after `EXTRAIL <extract-name>`, ensure that you include:

```
-- Table list for capture
table SRCMIRROR_OCIGLL.*;

-- Exclude changes made by GGADMIN
tranlogoptions excludeuser ggadmin
```

2. Add and run a Replicat to Database A. On the Parameters page, replace `MAP *.*`, `TARGET *.*`; with:

```
MAP SRCMIRROR_OCIGLL.*, TARGET SRC_OCIGLL.*;
```

3. Perform some changes on Database B to see them replicated to Database A.

Task 8: Monitor and maintain processes

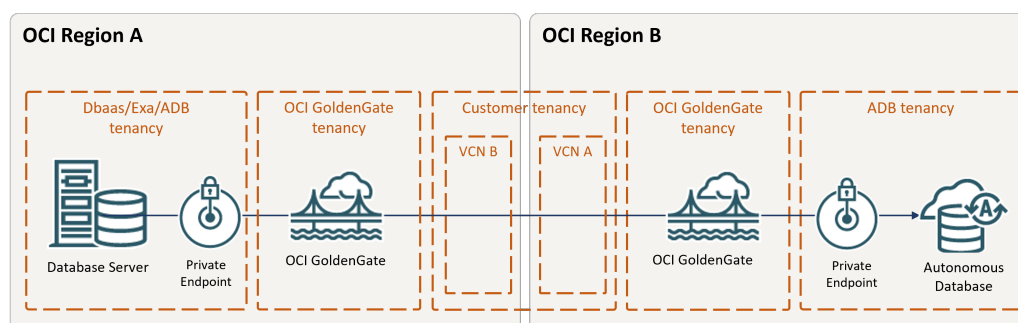
1. Monitor performance.
2. Manage Trail files.

Replicate data between cloud databases in different regions

Learn to set up and configure Oracle Cloud Infrastructure GoldenGate to replicate data between two Autonomous Databases located in different regions.

Overview

Oracle Cloud Infrastructure GoldenGate enables you to replicate data in supported OCI databases located in different regions. The following steps guide you through how to set up and run this replication model.



Before you begin

You must have the following in order to complete this quickstart:

- An existing source database in one region (Region A)
- An existing target database in different region (Region B)

See what's supported, and [in which regions OCI GoldenGate is available](#).

Task 1: Create the OCI GoldenGate resources

1. Create deployments in Regions A and B. Ensure that you enable **Create Public Endpoint** in the Advanced Options.
2. In Region A (source region), create a connection to the source database.
3. Assign the source connection to the source deployment in Region A.
4. In Region B (target region), create a connection to the target database.
5. Assign the target connection to the target deployment in Region B.

Task 2: Configure the source deployment

1. In Region A (source region), launch the deployment console from the deployment details page, and log in with the GoldenGate credentials you specified in Task 1.
2. Add Transaction information.
3. Add and run an Extract. Ensure that the Extract is running and capturing source changes before proceeding to the next step.
4. Add a credential that the target OCI GoldenGate deployment can use to connect to the source deployment:
 - a. Open the navigation menu and then click **Administrator**.
 - b. Click **Add User**, give the user a name (`ggsnet`, for example), and then assign the user the **Operator** role.

Task 3: Configure the target deployment

1. In Region B (target region), launch the deployment console from the deployment details page, and then log in using the GoldenGate credentials you specified in Task 1.
2. Add the source GoldenGate credential:
 - a. In the navigation menu, click **Configuration**.

- b. On the Credentials page, click **Add Credential**, and then complete the fields to add the source OCI GoldenGate **ggsnet** user from Task 2.
 - c. Click **Submit**.
3. Add and run a Receiver Path with the following values:
 - a. Source Authentication Method: **UserID Alias**
 - b. Source Protocol: **wss**
 - c. Source Host: `<domain>.deployment.goldengate.<source-region>-1.oci.oraclecloud.com`

 **Note:**

You can copy and paste the Console URL from the source Deployment Details page and remove the `https://` protocol.

- d. Source Port Number: 443
 - e. Source Trail Name: Enter the two character source trail name used when you created the Extract
 - f. Source Domain: Enter the source OCI GoldenGate user name (`ggsnet`)
 - g. Source Alias: Enter the source OCI GoldenGate alias
 - h. Target Trail: Enter a two character trail name for the target trail
4. Verify that the Receiver Path is created in the target region OCI GoldenGate Deployment Console.

Task 4: Replicate data

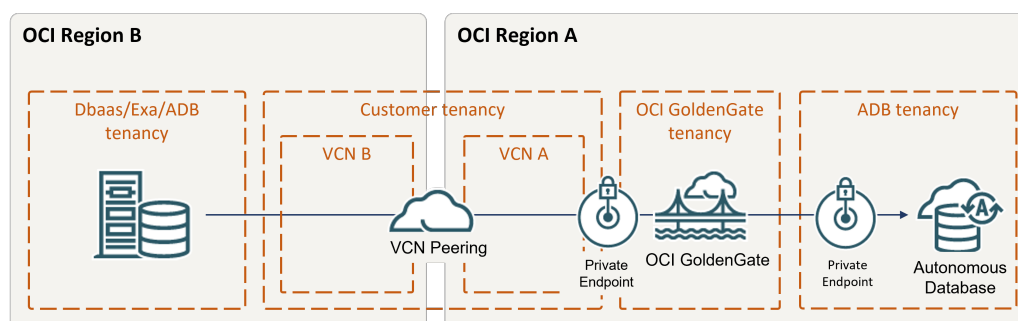
1. On the source deployment console in Region A, verify that the Distribution path was created in the Distribution Service.
2. Return the target deployment console in Region B, and then add and run a Replicat. The Replicat reads the Trail file created by the Receiver path.
3. Monitor performance.
4. Manage Trail files.

Replicate data between cloud databases in different regions with VCN peering

Learn to set up and configure Oracle Cloud Infrastructure GoldenGate and VCN peering to replicate data between two Autonomous Databases located in two different regions.

Overview

Oracle Cloud Infrastructure GoldenGate enables you to replicate data in supported OCI databases located in different regions with private endpoints. This example demonstrates how to connect OCI GoldenGate in Phoenix (Region A) to an Autonomous Transaction Processing (ATP) instance in Frankfurt (Region B) with a private endpoint.



Before you begin

You must have the following in order to proceed:

- An existing source database in one region (Region A)
- An existing target database in a different region (Region B)

Task 1: Configure networking

1. In Region A, [create a VCN](#) (VCN A) with two regional subnets:
 - Public (10.0.0.0/24)
 - Private (10.0.1.0/24)
 - a. On the **VCN A Details** page, under **Resources**, click **Dynamic Routing Gateway Attachments**, and then click **Create DRG Attachment**.
 - b. In the Create DRG Attachment panel, select the DRG you created, and then click **Create DRG Attachment**.
 - c. In the DRG Attachments list, click the DRG name in the Dynamic Routing Gateway column. You're brought to the DRG Details page.
 - d. On the DRG Details page, under **Resources**, click **Remote Peering Connection Attachments**, and then click **Create Remote Peering Connection**.
 - e. In the Create Remote Peering Connection panel, enter a name, leave the default settings as is, and then click **Create Remote Peering Connection**. An RPC attachment is automatically added to the DRG and its peering status set to New (not peered).
 - f. In the Remote Peering Connections Attachments list, under **Remote Peering Connection**, click the RPC name.
 - g. On the RPC Details page, for OCID, click **Copy**.

Note:

You can temporarily paste the OCID to a text editor for later use.

2. Repeat the previous step in Region B to [create a VCN](#) (VCN B) with two regional subnets and DRG:
 - Public (192.168.0.0/24)

- Private (1962.168.1.0/24)
3. On **Region B's RPC Details** page, click **Establish Connection**, select Region A's RPC, and then paste Region A's RPC OCID. The Peer Status is then set to Peered.
 4. On VCN A's Details page, under **Resources**, click **Route Tables**, and then click **Default Route Table for <VCN Name>**.
 5. Click **Add Route Rules**.
 6. In the Add Route Rules panel, complete the following fields, and then click **Add Route Rules**:
 - a. Target Type: **Dynamic Routing Gateway**
 - b. Destination CIDR Block: 192.168.0.0/24
 7. On VCN B's Details page, under **Resources**, click **Security Lists**, and then click **Default Security List for <VCN Name>**.
 8. Click **Add Ingress Rules**.
 9. In the Add Ingress Rules dialog, complete the following fields and then click **Add Ingress Rules**:
 - a. Source Type: **CIDR**
 - b. Source CIDR: 10.0.0.0/24
 - c. IP Protocol: **TCP**
 - d. Source Port Range: **All**
 - e. Destination Port Range: 1522

 **Note:**

This is the default port to access Oracle Autonomous Database (ADB) instances.

10. On VCN B's Details page, under **Resources**, click **Route Tables**, and then click **Default Route Table for <VCN Name>**.
11. Click **Add Route Rules**.
12. In the Add Route Rules panel, complete the following fields and then click **Add Route Rules**:
 - a. Target Type: Dynamic Routing Gateway
 - b. Destination CIDR: 10.0.0.0/24

Task 2: Create a deployment

Ensure that you use VCN A in Region A, which was peered with VCN B in Region B.

To see which regions OCI GoldenGate is available in, see [Cloud Data Regions](#).

1. In the Console navigation menu, click **Oracle Database**, and then select **GoldenGate**.
2. On the Deployments page, click **Create deployment**.
3. In the Create deployment panel, enter a name and optionally, a description.

4. From the **Compartment** dropdown, select a compartment in which to create the deployment.
5. Select one of the following options:
 - **Production:** Sets up a deployment with recommended defaults for a production environment. The minimum number of OCPUs is 4, with auto-scaling enabled.
 - **Development or testing:** Sets up a deployment with recommended defaults for a development or testing environment. The minimum number of OCPUs is 1.
6. For **OCPU count** enter the number of Oracle Compute units (OCPUs) to use.

 **Note:**

One OCPU is equivalent to 16gb of memory. For more information, see [OCPU management and billing](#).

7. (Optional) Select **Auto scaling**.

 **Note:**

Auto scaling enables OCI GoldenGate to scale up to three times the number of OCPUs you specify for OCPU Count, up to 24 OCPUs. For example, if you specify your OCPU Count as 2 and enable Auto Scaling, then your deployment can scale up to 6 OCPUs. If you specify your OCPU Count as 20 and enable Auto Scaling, OCI GoldenGate can only scale up to 24 OCPUs.

8. From the **Subnet in <Compartment>** dropdown, select the subnet to which a private endpoint is created from the OCI GoldenGate service tenancy. This ensures that the deployment is always available over this subnet, as long as the policies for this subnet allow access. The private endpoint is only used to access the deployment console, and doesn't provide access to other resources in the subnet.

To select a subnet in a different compartment, click **Change compartment**.

 **Note:**

You can only select a private subnet when creating a deployment.

9. Select a license type.
10. (Optional) Click **Show advanced options** for network options and to add tags.
 - a. In the **Network** tab,
 - i. Select **Enable GoldenGate console public access** to include a public endpoint in addition to a private endpoint, and allow public access to the deployment console for users. If selected, OCI GoldenGate creates a load balancer in your tenancy to create a public IP. Select a subnet in the same VCN as this deployment in which to create the load balancer.

 **Note:**

The load balancer is a resource that comes with an additional cost. You can manage this resource, but ensure that you don't delete the load balancer while your deployment is still in use. [Learn more about load balancer pricing.](#)

- ii. Select **Customize endpoint** to provide a private fully qualified domain name (FQDN) prefix that you'll use to access the private service console URL. You can also optionally upload an SSL/TLS certificate (.pem) and its corresponding private key, however, password protected certificates are not supported. It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.
A self-signed certificate is generated for you, if you don't provide one.

 **Note:**

It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.

- b. In the **Maintenance** tab:
 - i. Select **Customize maintenance window** to define the start of the maintenance window to upgrade the deployment.
 - ii. (Optional) For **Major release auto-upgrade period in days**, enter the number of days, between 0 and 365.
 - iii. (Optional) For **Bundle release auto-upgrade period in days**, enter the number of days, between 0 and 180 days.
 - iv. (Optional) For **Security patch auto-upgrade period in days**, enter the number of days, between 0 and 14 days.
 - v. Select **Enable interim release auto-upgrade**, and, optionally, enter the number of days.

 **Note:**

[Learn more about scheduling upgrades.](#)

- c. In the **Tags** tab, add tags to help track the resources within your tenancy. Click **+ Additional tag** to add more tags. [Learn more about tagging.](#)
11. Click **Next**.
 12. For Deployment type, select **Data replication**.
 13. From the **Select a technology** dropdown, select one of the following technology types:
 - Oracle Database
 - Big Data
 - MySQL
 - PostgreSQL

- Microsoft SQL Server

See what's supported to learn which databases and technologies you can use as OCI GoldenGate sources and targets.

14. For **Version**, the latest version is automatically selected. Click **Change version** to select a different version.

 **Note:**

Learn more about versions.

15. For **GoldenGate instance name**, enter the name that the deployment will assign to the GoldenGate deployment instance upon creation.
16. For Credential store, select one of the following:
 - **OCI Identity and Access Management (OCI IAM)**, to enable users to log in to the the deployment console using their Oracle Cloud account (single sign on) in IAM (Identity and Access Management) enabled tenancies.

 **Note:**

Once you select IAM, you won't be able to switch to GoldenGate when you edit the deployment settings at a later time.

- **GoldenGate**, for GoldenGate to manage users.
 - a. Enter the **Administrator username**
 - b. Select a password secret in your compartment or click **Change compartment** to select one in a different compartment. You can also create a new password secret.

To create a new password secret:

 - i. Click **Create password secret**.
 - ii. In the Create secret panel, enter a name for the secret, and optionally, a description.
 - iii. Select a compartment from the **Compartment** dropdown in which to save your secret.
 - iv. Select a vault in the current compartment, or click **Change compartment** to select a vault in a different compartment.
 - v. Select an **Encryption key**.

 **Note:**

Only AES keys, Software protected keys, and HSM keys are supported. RSA and ECDSA keys are not supported for GoldenGate password secret keys.

- vi. Enter a password 8 to 30 characters in length, containing at least 1 uppercase, 1 lowercase, 1 numeric and 1 special character. The special characters must not be '\$', '^' or '?'.

- vii. Confirm the password.
- viii. Click **Create**.

 **Note:**

You can manage GoldenGate users in the deployment console. [Learn more.](#)

17. Click **Create**.

Task 3: Create and assign connections

1. Create connections for the source and target databases.
2. Assign the connections to the deployment created in Task 2.

Task 4: Replicate data

1. Navigate back to the Deployments page, and then select the deployment you created in Task 2.
2. On the Deployment details page, click **Launch console**.
3. Log in to the OCI GoldenGate deployment console
4. Add transaction information and a checkpoint table.
5. Add and run an Extract.
6. Add and run a Replicat.

Task 5: Monitor and maintain processes

1. Monitor the replication process.
2. Manage Trail files.

Learn more

- [Remote VCN Peering using an RPC](#)

Send data from Oracle GoldenGate to OCI GoldenGate

Learn to create a trusted connection and send data from an on-premises or Marketplace Oracle GoldenGate to Oracle Cloud Infrastructure GoldenGate

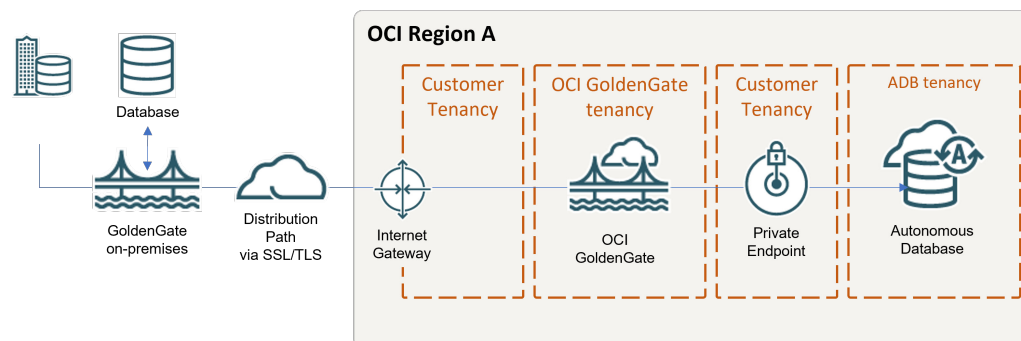
 **Note:**

This quickstart applies only to Oracle GoldenGate Microservices because OCI GoldenGate only allows SSL-based communication. If you're looking for how to replicate data from Oracle GoldenGate Classic to OCI GoldenGate, see [Connecting GoldenGate Classic to GoldenGate Microservices and OCI GoldenGate](#).

Overview

Oracle Cloud Infrastructure GoldenGate enables you to send data from an on-premises or Marketplace Oracle GoldenGate to OCI GoldenGate using a Distribution Path. The following steps guide you through how to set up and run this replication model using the latest Oracle GoldenGate version for OCI Marketplace.

This quickstart is also available as a LiveLab: [View the workshop](#).



Before you begin

You must have the following in order to proceed:

- An existing on-premises or Marketplace Oracle GoldenGate deployment
- An existing source and target Autonomous databases.

Note:

[Download](#) and load the sample data, if needed.

Task 1: Create OCI GoldenGate resources

1. Create an OCI GoldenGate deployment.
2. Create connections to your source and target databases.
3. Assign connections to the deployment.

Task 2: Create a trusted connection between Oracle GoldenGate and OCI GoldenGate

Take care in distinguishing the Oracle GoldenGate Service Manager from the OCI GoldenGate deployment console while you complete this task.

1. Download the root certificate for the **OCI GoldenGate Deployment Console**.

Note:

You can download the root certificate from any browser. The following steps describe how to download the root certificate from a Chrome browser.

- a. In your Chrome browser address bar, click the padlock icon, and then click **Connection is secure**.
 - b. Click **Certificate is valid**. A Certificate window opens.
 - c. In the Certificate window, click **Certification Path**, select **DigiCert**, and then click **View Certificate**.
 - d. Ensure that **Issued** reads **DigiCert Global Root G2**, click **Details**, and then **Copy to File**.
 - e. In the Certificate Export Wizard, click **Next**, select **Base-64 encoded X.509 (.CER)**, and then click **Next**.
 - f. Click **Browse** to save the file to your local drive.
 - g. Click **Next**, and then click **Finish**.
2. Upload the certificate to the on-premises or Marketplace Oracle GoldenGate Service Manager:
 - a. Open the on-premises or Marketplace Oracle GoldenGate Service Manager in a browser window.
 - b. In the navigation menu (hamburger icon), click **Certificate Management**.
 - c. On the Certificate Management page, click **Add CA Certificates** (plus icon).
 - d. In the Add CA Certificate dialog, enter a **Unique Name**.
 - e. For **Certificate PEM**, copy and paste the contents of the certificate file you downloaded from Chrome, and then click **Add**.
 3. Add a Credential that allows Oracle GoldenGate to connect to OCI GoldenGate.
 - a. Launch the OCI GoldenGate Deployment Console from the Deployment Details page.
 - b. Log in, and then access the **Administrator** page from the navigation menu.
 - c. Click **Add User**, and then create a user that Oracle GoldenGate can use to connect to OCI GoldenGate. Assign this user the **Operator** role.
 - d. Open the on-premises or Marketplace Oracle GoldenGate Administration Service, and then navigate to the Configuration page.
 - e. Under the Database tab, click **Add Credential**, and then complete the following:
 - **Credential Domain**: Enter a name for this connection
 - **Credential Alias**: Enter an alias
 - **User ID**: Enter the name of the user created in step 5c.
 - **Password** and **Verify Password**: Enter the password associated with this user.
 - f. Click **Submit**.

Task 3: Send data from Oracle GoldenGate to OCI GoldenGate

1. On the on-premises or Marketplace Oracle GoldenGate, add and run an extract.
2. On the on-premises or Marketplace Oracle GoldenGate, add a distribution path with the following values, and then click **Create and Run**:
 - a. **Path Name**: Enter a name for this path
 - b. **Source**: Select the Extract created in step 1.

- c. **Trail file:** Select the trail file to send to OCI GoldenGate
- d. **Target Authentication Method:** UserID Alias
- e. **Target Host:** Enter the OCI GoldenGate hostname in the following format, `<domain>.deployment.goldengate.<region>.oci.oraclecloud.com:443`

 **Note:**

You can copy and paste the Console URL From your OCI GoldenGate deployment details page, and remove the `https://` protocol and any trailing slashes (`/`).

- f. **Target Trail Name:** Enter a two-character name for the trail when it's received by OCI GoldenGate
 - g. **Target Domain:** Enter the domain name you created in Task 2.
 - h. **Target Alias:** Enter the alias name you created in Task 2.
3. In the OCI GoldenGate deployment console, add a non-integrated replicat.

Task 4: Monitor performance

- 1. Monitor performance.
- 2. Manage Trail files.

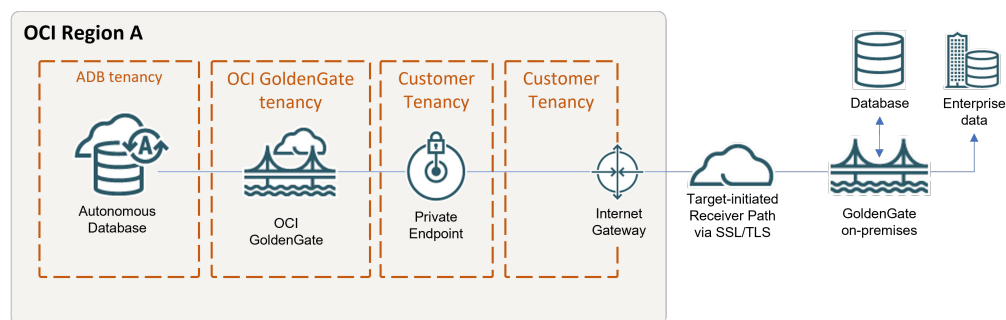
Send data from OCI GoldenGate to Oracle GoldenGate

Learn to create a trusted connection and pull data from Oracle Cloud Infrastructure GoldenGate to Oracle GoldenGate.

Overview

Oracle Cloud Infrastructure GoldenGate enables you to send data from an OCI GoldenGate deployment to an on-premises or Marketplace Oracle GoldenGate deployment using a Receiver Path. The following steps guide you through how to set up and run this replication model using the latest Oracle GoldenGate Marketplace instance.

This quickstart is also available as a LiveLab: [View the workshop](#).



Before you begin

You must have the following in order to proceed:

- An existing source database.
- An existing on-premises or Marketplace Oracle GoldenGate deployment

Task 1: Create OCI GoldenGate resources

1. Create an OCI GoldenGate deployment.
2. Create connections to your source and target databases.
3. Assign connections to the deployment.

Task 2: Create a trusted connection between Oracle GoldenGate and OCI GoldenGate

Take care in distinguishing the Oracle GoldenGate Service Manager from the OCI GoldenGate deployment console while you complete this task.

1. Download the root certificate for the **OCI GoldenGate Deployment Console**.

 **Note:**


You can download the root certificate from any browser. The following steps describe how to download the root certificate from a Chrome browser.

- a. In your Chrome browser address bar, click the padlock icon, and then click **Connection is secure**.
 - b. Click **Certificate is valid**. A Certificate window opens.
 - c. In the Certificate window, click **Certification Path**, select **DigiCert**, and then click **View Certificate**.
 - d. Ensure that **Issued** reads **DigiCert Global Root G2**, click **Details**, and then **Copy to File**.
 - e. In the Certificate Export Wizard, click **Next**, select **Base-64 encoded X.509 (.CER)**, and then click **Next**.
 - f. Click **Browse** to save the file to your local drive.
 - g. Click **Next**, and then click **Finish**.
2. Upload the certificate to the on-premises or Marketplace Oracle GoldenGate Service Manager:
 - a. Open the on-premises or Marketplace Oracle GoldenGate Service Manager in a browser window.
 - b. In the navigation menu (hamburger icon), click **Certificate Management**.
 - c. On the Certificate Management page, click **Add CA Certificates** (plus icon).
 - d. In the Add CA Certificate dialog, enter a **Unique Name**.
 - e. For **Certificate PEM**, copy and paste the contents of the certificate file you downloaded from Chrome, and then click **Add**.
 3. Add a Credential that allows Oracle GoldenGate to connect to OCI GoldenGate.
 - a. Launch the OCI GoldenGate Deployment Console from the Deployment Details page.

- b. Log in, and then access the **Administrator** page from the navigation menu.
- c. Click **Add User**, and then create a user that Oracle GoldenGate can use to connect to OCI GoldenGate. Assign this user the **Operator** role.
- d. Open the on-premises or Marketplace Oracle GoldenGate Administration Service, and then navigate to the Configuration page.
- e. Under the Database tab, click **Add Credential**, and then complete the following:
 - **Credential Domain**: Enter a name for this connection
 - **Credential Alias**: Enter an alias
 - **User ID**: Enter the name of the user created in step 5c.
 - **Password** and **Verify Password**: Enter the password associated with this user.
- f. Click **Submit**.

Task 3: Send data from OCI GoldenGate to Oracle GoldenGate

This task instructs you on how to create and run a target-initiated Receiver Path to pull Trail files from OCI GoldenGate to Oracle GoldenGate.

1. In the OCI GoldenGate deployment console Oracle GoldenGate, add and run an extract.
 2. On the on-premises or Marketplace Oracle GoldenGate, add a receiver path with the following values, and then click **Create and Run**:
 - a. **Path Name**: Enter a name for this path
 - b. **Source**: Select the Extract created in step 1.
 - c. **Source Authentication Method**: Use Basic Authentication to UserID Alias
 - d. **Source Host**: Enter the OCI GoldenGate hostname in the following format, `<domain>.deployment.goldengate.<region>.oci.oraclecloud.com:443`
-  **Note:**

You can copy and paste the Console URL From your OCI GoldenGate deployment details page, and remove the `https://` protocol and any trailing slashes (`/`).
- e. **Source Trail Name**: Enter the name of the OCI GoldenGate trail file to send to Oracle GoldenGate
 - f. **Source Domain**: Enter the domain name you created in Task 2.
 - g. **Source Alias**: Enter the alias name you created in Task 2.
 - h. **Target**: Enter a two-character name for the Trail file when it is received by Oracle GoldenGate
3. On the on-premises or Marketplace Oracle GoldenGate, add a replicat.

Task 4: Monitor performance

1. Monitor performance.

2. Manage Trail files.

Big Data quickstarts

Common use cases using Big Data technologies as OCI GoldenGate sources or targets.

Articles in this section:

- [Replicate data from Autonomous Transaction Processing to OCI Object Storage](#)
- [Replicate data from Amazon RDS for Oracle to OCI Object Storage](#)
- [Replicate data from Autonomous Database to OCI Streaming](#)
- [Replicate data from Autonomous Transaction Processing to Apache Kafka](#)
- [Capture data from Kafka platforms](#)
- [Stage and merge data into Autonomous Data Warehouse using OCI GoldenGate](#)
- [Replicate Data from Autonomous Transaction Processing to Confluent Kafka](#)
- [Replicate data from Autonomous Transaction Processing to Azure Data Lake Storage Gen 2](#)
- [Replicate Data from Autonomous Transaction Processing to Azure Synapse Analytics](#)
- [Replicate data from Autonomous Transaction Processing to Amazon S3](#)
- [Replicate data from MongoDB to Autonomous JSON Database](#)
- [Replicate data from PostgreSQL to Snowflake](#)
- [Replicate data from MySQL HeatWave to Amazon Kinesis](#)
- [Send data from MySQL HeatWave to Azure Event Hubs](#)
- [Replicate data from MySQL HeatWave to Google Cloud Storage](#)
- [Replicate Data from PostgreSQL to Google BigQuery](#)

Replicate data from Autonomous Database to OCI Object Storage

This quickstart example demonstrates how to replicate data from Autonomous Transaction Processing to OCI Object Storage using OCI GoldenGate.

This quickstart is also available as a LiveLab workshop. [View the workshop](#).

Before you begin

To successfully complete this quickstart, you must have the following:

- An [API key created and downloaded](#) from the user details page
- A source Autonomous Transaction Processing instance
- Your [compartment OCID](#)

Task 0: Set up the source Autonomous Database

If you don't already have a source database set up for replication, you can follow these steps to load a sample schema to use for this quickstart. This quickstart uses Autonomous Transaction Processing (ATP) for the source database.

To set up the source Autonomous Database:

1. Download and unzip the [sample database schema](#).
2. In the Oracle Cloud console, select your ATP instance from the Autonomous Databases page to view its details and access Database Actions.
3. Click **Open Database Actions**.
4. Enable the GGADMIN user:
 - a. Under **Administration**, click **Database Users**.
 - b. Locate **GGADMIN** and then click its ellipsis menu (three dots) and select **Edit**.
 - c. In the Edit User panel, enter the GGADMIN password, confirm the password, and then disable **Account is Locked**.
 - d. Click **Apply Changes**.
5. Load the load the source sample schema and data:
 - a. From the Database Actions Selector menu, under **Development**, select **SQL**.
 - b. Copy and paste the script from **OCIGLL_OCIGGS_SETUP_USERS_ATP.sql** into the SQL worksheet.
 - c. Click **Run Script**. The Script Output tab displays confirmation messages.
 - d. Clear the SQL worksheet and then copy and paste the SQL script from **OCIGLL_OCIGGS_SRC_USER_SEED_DATA.sql**.

 **Tip:**

You may need to run each statement separately for the SQL tool to execute the scripts successfully.

- e. To verify that the tables were created successfully, close the SQL window and reopen it again. In the Navigator tab, look for the SRC_OCIGLL schema and then select tables from their respective dropdowns.
6. Enable supplemental logging:
 - a. Clear the SQL Worksheet.
 - b. Enter the following statement, and then click **Run Statement**:

```
ALTER PLUGGABLE DATABASE ADD SUPPLEMENTAL LOG DATA;
```

Task 1: Create the OCI GoldenGate resources

This quickstart example requires deployments and connections for both the source and target.

1. Create an Oracle deployment for the source Autonomous Transaction Processing instance.
2. Create a Big Data deployment for the target OCI Object Storage bucket.
3. Create a connection for the source Autonomous Transaction Processing instance.
4. Create connection for the target OCI Object Storage.

5. (Optional) If your Big Data deployment doesn't have a public endpoint, then create a connection to GoldenGate, and then assign this connection to the source Oracle deployment.
6. Assign the Autonomous Transaction Processing connection to the source Oracle deployment.
7. Assign the OCI Object Storage connection to the target Big Data deployment.

Task 2: Add and run the Extract

1. On the Deployments page, select the source Autonomous Transaction Processing deployment.
2. On the deployment details page, click **Launch Console**.
3. Log in with the source deployment's administrator username and password.
4. Add transaction information.
5. Add an Extract.

Task 3: Add and run a Distribution Path

1. Create a user for the Distribution Path in the target Big Data deployment.
 - a. On the Deployments page, select the target deployment to view its details.
 - b. On the deployment details page, click **Launch Console**.
 - c. Sign in to the target OCI GoldenGate deployment console using the Administrator username and password.
 - d. In the OCI GoldenGate deployment console, open the navigation menu, and then click **Administrator**.
 - e. On the Users page, click **Add New User** (plus icon).
 - f. Complete the fields as follows, and then click **Submit**.
 - For **Username**, enter a name, such as `ggsnet`.
 - From the **Role** dropdown, select **Operator**.
 - Enter a password twice for verification.

The new user appears in the Users list.

2. In the source ATP GoldenGate deployment console, add a credential for the user created in Step 1.
 - a. In the source ATP GoldenGate deployment console, click **Administration Service**, open the navigation menu, and then select **Configuration**.
 - b. On the Credentials page, click **Add Credential**, and then complete the fields as follows:
 - For **Credential Domain**, enter `GGNetwork`.
 - For **Credential Alias**, enter `dpuser`
 - For **User ID**, enter the name of the user created in Step 1 (`ggsnet`)
 - Enter the user's password twice for verification.

- c. Click **Submit**.The credential appears in the Credentials list.
3. Click **Distribution Service**, and then click **Add Path** (plus icon).
4. Complete the Add Path form fields as follows, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.
 - b. For **Source Extract**, select the Extract created in Task 2.
 - c. For **Source Trail Name**, select the Extract Trail from Task 2.
 - d. For **Target Authentication Method**, select **UserID Alias**.
 - e. For **Target** protocol, select **wss**.
 - f. For **Target Host**, enter the host domain of the target deployment.

 **Note:**

You can copy and paste the URL of the target OCI GoldenGate deployment console and remove the `https://` and any trailing slashes.

- g. For **Port Number**, enter 443.
 - h. For **Trail Name**, enter a two-character name, such as E1.
 - i. For **Domain**, enter the name of the Credential Domain created in Step 2 (GGSNetwork).
 - j. For **Alias**, enter the Credential Alias created in Step 2 (dpuser).

You're returned to the Distribution Service Overview page where you can view the status of the created path.

5. In the target Big Data deployment console, review the Receiver Path.
 - a. In the target Big Data OCI GoldenGate deployment console, click **Receiver Service**.
 - b. Review the path details. This path was created as a result of the Distribution Path created in the previous step.

Task 4: Add and run the Replicat

1. In the target Big Data deployment console, click **Administrator Service**, and then click **Add Replicat** (plus icon).
2. On the Add Replicat page, under Replicat type, select **Classic Replicat**, and then click **Next**.
3. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name.
 - For **Trail Name**, enter the name of the Trail from Task 2.
 - For **Target**, select **OCI Object Storage**.
 - For **Alias**, select the OCI Object Storage connection created in Task 1.

4. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP SRC_OCIGLL.*, TARGET *.*;
```

5. On the Properties page, configure OCI Event Handler properties as needed, and then click **Next**.

Required properties:

- `gg.eventhandler.oci.compartmentID`: The compartment OCID in which the OCI Object Storage bucket resides.
- `gg.handler.oci.fileNameMappingTemplate`: Generates files names dynamically using Template Keywords.

Properties you may consider modifying include:

- `gg.handler.oci.format`: Select how to format the output. `json_row` is the default setting. Available options include:
 - `delimitedtext`
 - `json`
 - `json_row`
 - `xml`
 - `avro_row_ocf`
 - `avro_op_ocf`

**Tip:**

To use the formatting property for OCI Object Storage, replace `name` with `oci`. For example, `gg.handler.name.format` becomes `gg.handler.oci.format`.

- `gg.handler.oci.inactivityRollInterval`: GoldenGate creates a file and keeps it open for writing. This property closes the file after the designated period of inactivity (no incoming transactions), and then loads it into OCI Object Storage. By default, it is set to 5 seconds. You can specify a time in milliseconds (ms), seconds (s), minutes (m), or hours (h). For example, `gg.handler.oci.inactivityRollInterval=10m`.
- `gg.handler.oci.maxFileSize`: File Writer Handler opens the file and keeps it open until it reaches the maximum file size, assuming there are no metadata changes. By default, the maximum file size is 1 GB, however you can change it using this property. When the size is reached, file is closed, and a new file is generated. For example, `gg.handler.oci.maxFileSize=500m`.
- `gg.handler.oci.rollOnShutdown`: The default value is `true`. When set to `true`, GoldenGate shuts the open file when you stop the Replicat process. By default, File Writer Handler keeps the file open even if the Replicat stops and continues writing to the same file when the Replicat restarts. For example, `gg.handler.oci.rollOnShutdown=false`
- `gg.handler.oci.fileRollInterval`: Designates the amount of time to keep the file open before it's closed and rolls over to a new file. By default it is set to 7 minutes. You can specify a time in milliseconds (ms), seconds (s), minutes (m), or hours (h). For example, `gg.handler.oci.fileRollInterval=10m`.

- `gg.eventhandler.oci.bucketMappingTemplate`: Enter the Object Storage bucket name.
6. Click **Create and Run**.
You're returned to the Overview page, where you can review the Replicat details.

Task 5: Verify the replication

Perform some updates to the source Autonomous Transaction Processing instance to verify replication to OCI Object Storage.

1. In the Oracle Cloud console, open the navigation menu, select **Oracle Database**, and then select **Autonomous Transaction Processing**.
2. In the list of Autonomous Transaction Processing instances, select your source instance to view its details.
3. On the database details page, click **Database Actions**.

Note:

You should be automatically logged in. If not, log in with the database credentials.

4. On the Database Actions home page, select **SQL**.
5. Enter the following into the worksheet and click **Run Script**.

```
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1000,'Houston',20,743113);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1001,'Dallas',20,822416);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1002,'San
Francisco',21,157574);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1003,'Los
Angeles',21,743878);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1004,'San
Diego',21,840689);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1005,'Chicago',23,616472);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1006,'Memphis',23,580075);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1007,'New York
City',22,124434);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1008,'Boston',22,275581);
```

```
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1009,'Washington D.C.',22,688002);
```

6. In the source GoldenGate OCI GoldenGate deployment console, select the Extract name, and then click **Statistics**. Verify that **SRC_OCIGLL.SRC_CITY** has 10 inserts.
7. In the target Big Data OCI GoldenGate deployment console, select the Replicat name, and then click **Statistics**. Verify that **SRC_OCIGLL.SRC_CITY** has 10 inserts.
8. In the Oracle Cloud console, navigate to the OCI Object Storage bucket and check its contents.

Task 6: Monitor and maintain processes

1. Monitor performance.
2. Manage Trail files.

Replicate data from Autonomous Database to OCI Streaming

Learn to replicate data from a source Autonomous Transaction Processing instance to OCI Streaming using OCI GoldenGate.

Before you begin

To successfully complete this quickstart, you must have the following:

- [An Auth Token created](#) and downloaded from the user details page
- [A stream created](#) in the OCI Streaming service
- A source Autonomous Transaction Processing instance

Task 0: Set up the source Autonomous Database

If you don't already have a source database set up for replication, you can follow these steps to load a sample schema to use for this quickstart. This quickstart uses Autonomous Transaction Processing (ATP) for the source database.

To set up the source Autonomous Database:

1. Download and unzip the [sample database schema](#).
2. In the Oracle Cloud console, select your ATP instance from the Autonomous Databases page to view its details and access Database Actions.
3. Click **Open Database Actions**.
4. Enable the GGADMIN user:
 - a. Under **Administration**, click **Database Users**.
 - b. Locate **GGADMIN** and then click its ellipsis menu (three dots) and select **Edit**.
 - c. In the Edit User panel, enter the GGADMIN password, confirm the password, and then disable **Account is Locked**.
 - d. Click **Apply Changes**.
5. Load the load the source sample schema and data:
 - a. From the Database Actions Selector menu, under **Development**, select **SQL**.

- b. Copy and paste the script from **OCIGLL_OCIGGS_SETUP_USERS_ATP.sql** into the SQL worksheet.
- c. Click **Run Script**. The Script Output tab displays confirmation messages.
- d. Clear the SQL worksheet and then copy and paste the SQL script from **OCIGLL_OCIGGS_SRC_USER_SEED_DATA.sql**.

 **Tip:**

You may need to run each statement separately for the SQL tool to execute the scripts successfully.

- e. To verify that the tables were created successfully, close the SQL window and reopen it again. In the Navigator tab, look for the SRC_OCIGLL schema and then select tables from their respective dropdowns.
6. Enable supplemental logging:
 - a. Clear the SQL Worksheet.
 - b. Enter the following statement, and then click **Run Statement**:

```
ALTER PLUGGABLE DATABASE ADD SUPPLEMENTAL LOG DATA;
```

Task 1: Create the OCI GoldenGate resources

This quickstart example requires deployments and connections for both the source and target.

1. Create an Oracle deployment for the source Autonomous Transaction Processing instance.
2. Create a Big Data deployment for the target Stream.
3. Create a connection for the source Autonomous Transaction Processing instance.
4. Create a connection to OCI Streaming.
5. (Optional) If your Big Data deployment doesn't have a public endpoint, then create a connection to GoldenGate, and then assign this connection to the source Oracle deployment.
6. Assign the ATP connection to the source Oracle deployment.
7. Assign the OCI Streaming connection to the target Big Data deployment.

Task 2: Add the Extract

1. On the Deployments page, select the source Autonomous Transaction Processing deployment.
2. On the deployment details page, click **Launch Console**.
3. Log in with the source deployment's administrator username and password.
4. Add transaction information.
5. Add an Extract

Task 3: Add a Distribution Path

1. Create a user for the Distribution Path in the target Big Data deployment console.
 - a. On the Deployments page, select the target deployment to view its details.
 - b. On the deployment details page, click **Launch Console**.
 - c. Sign in to the target OCI GoldenGate deployment console using the Administrator username and password.
 - d. In the OCI GoldenGate deployment console, open the navigation menu, and then click **Administrator**.
 - e. On the Users page, click **Add New User** (plus icon).
 - f. Complete the fields as follows, and then click **Submit**.
 - For **Username**, enter a name, such as `ggsnet`.
 - From the **Role** dropdown, select **Operator**.
 - Enter a password twice for verification.

The new user appears in the Users list.

2. In the source ATP GoldenGate deployment, add a credential for the user created in Step 1.
 - a. In the source ATP deployment console, click **Administration Service**, open the navigation menu, and then select **Configuration**.
 - b. On the Credentials page, click **Add Credential**, and then complete the fields as follows:
 - For **Credential Domain**, enter `GGSNetwork`.
 - For **Credential Alias**, enter `dpuser`
 - For **User ID**, enter the name of the user created in Step 1 (`ggsnet`)
 - Enter the user's password twice for verification.
 - c. Click **Submit**.

The credential appears in the Credentials list.

3. Click **Distribution Service**, and then click **Add Path** (plus icon).
4. Complete the Add Path form fields as follows, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.
 - b. For **Source Extract**, select the Extract created in Task 2.
 - c. For **Source Trail Name**, select the Extract Trail from Task 2.
 - d. For **Target Authentication Method**, select **UserID Alias**.
 - e. For **Target** protocol, select **wss**.
 - f. For **Target Host**, enter the host domain of the target deployment.

 **Note:**

You can copy and paste the URL of the target OCI GoldenGate deployment console and remove the `https://` and any trailing slashes.

- g. For **Port Number**, enter 443.
 - h. For **Trail Name**, enter a two-character name, such as E1.
 - i. For **Domain**, enter the name of the Credential Domain created in Step 2 (GGSNetwork).
 - j. For **Alias**, enter the Credential Alias created in Step 2 (dpuser).
- You're returned to the Distribution Service Overview page where you can view the status of the created path.
- 5. In the target Big Data deployment console, review the Receiver Path.
 - a. In the target Big Data OCI GoldenGate deployment console, click **Receiver Service**.
 - b. Review the path details. This path was created as a result of the Distribution Path created in the previous step.

Task 4: Add the Replicat

1. In the target Big Data deployment console, click **Administrator Service**, and then click **Add Replicat** (plus icon).
2. On the Add Replicat page, under Replicat type, select **Classic Replicat**, and then click **Next**.
3. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name.
 - For **Trail Name**, enter the name of the Trail from Task 2.
 - For **Target**, select **OCI Streaming**.
 - For **Alias**, select the Stream connection created in Task 1.
4. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP SRC_OCIGLL.*, TARGET *.*;
```

5. On the Properties page, use the Stream name values to set the following property:


```
gg.handler.kafkahandler.topicMappingTemplate=<stream-name>
```
6. Click **Create and Run**.

You're returned to the Overview page, where you can review the Replicat details.

Task 5: Verify the replication

Perform some updates to the source Autonomous Transaction Processing instance to verify replication to OCI Streaming.

1. In the Oracle Cloud console, open the navigation menu, select **Oracle Database**, and then select **Autonomous Transaction Processing**.
2. In the list of Autonomous Transaction Processing instances, select your source instance to view its details.
3. On the database details page, click **Database Actions**.

 **Note:**

You should be automatically logged in. If not, log in with the database credentials.

4. On the Database Actions home page, select **SQL**.
5. Enter the following into the worksheet and click **Run Script**.

```
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1000,'Houston',20,743113);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1001,'Dallas',20,822416);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1002,'San Francisco',21,157574);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1003,'Los Angeles',21,743878);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1004,'San Diego',21,840689);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1005,'Chicago',23,616472);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1006,'Memphis',23,580075);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1007,'New York City',22,124434);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1008,'Boston',22,275581);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1009,'Washington D.C.',22,688002);
```

6. In the source ATP deployment console, select the Extract name, and then click **Statistics**. Verify that **SRC_OCIGLL.SRC_CITY** has 10 inserts.
7. In the target Big Data deployment console, select the Replicat name, and then click **Statistics**. Verify that **SRC_OCIGLL.SRC_CITY** has 10 inserts.
8. In the Oracle Cloud console, navigate to OCI Streaming and check the contents of your stream.

Task 6: Monitor and maintain processes

1. Monitor performance.
2. Manage Trail files.

Replicate data from Autonomous Transaction Processing to Amazon S3

Learn to replicate data from Autonomous Transaction Processing (ATP) to an Amazon S3 (Amazon Simple Storage Service) using OCI GoldenGate.

Before you begin

To successfully complete this quickstart, you must have the following:

- A source Autonomous Transaction Processing (ATP) instance
- Amazon S3 Service
- [Amazon S3 Access Key & Secret](#)

Set up the environment

If you don't already have a source database set up for replication, you can follow these steps to load a sample schema to use for this quickstart. This quickstart uses Autonomous Transaction Processing (ATP) for the source database.

To set up the source Autonomous Database:

1. Download and unzip the [sample database schema](#).
2. In the Oracle Cloud console, select your ATP instance from the Autonomous Databases page to view its details and access Database Actions.
3. On the Database details page, click **Database actions**.
4. Enable the GGADMIN user:
 - a. On the Database actions page, under Administration, click **Database Users**.
 - b. Locate GGADMIN and then click its ellipsis menu (three dots) and select **Edit**.
 - c. In the Edit User panel, enter the GGADMIN password, confirm the password, and then deselect **Account is Locked**.
 - d. Click **Apply Changes**.
5. Load the source sample schema and data:
 - a. From the Database Actions Selector menu, under Development, select **SQL**.
 - b. Copy and paste the script from `OCIGLL_OCIGGS_SETUP_USERS_ATP.sql` into the SQL worksheet.
 - c. Click **Run Script**. The Script Output tab displays confirmation messages.
 - d. Clear the SQL worksheet and then copy and paste the SQL script from `OCIGLL_OCIGGS_SRC_USER_SEED_DATA.sql`.

Tip:

You may need to run each statement separately for the SQL tool to execute the scripts successfully.

- e. To verify that the tables were created successfully, close the SQL window and reopen it again. In the Navigator tab, look for the `SRC_OCIGLL` schema and then select tables from their respective dropdowns.

6. Enable supplemental logging:
 - a. Clear the SQL Worksheet.
 - b. Enter the following statement, and then click **Run Statement**:

```
ALTER PLUGGABLE DATABASE ADD SUPPLEMENTAL LOG DATA;
```

Task 1: Create the OCI GoldenGate Resources

This quickstart example requires deployments and connections for both the source and target.

1. Create an Oracle deployment for the source Autonomous Transaction Processing instance.
2. Create a Big Data deployment for the target Amazon S3.
3. Create a connection for the source Autonomous Transaction Processing instance.
4. Create a connection to the target Amazon S3.
5. (Optional) If your Big Data deployment doesn't have a public endpoint, then create a connection to GoldenGate pointing to private end point of Big Data deployment, and then assign this connection to the source Oracle deployment.
6. Assign the Autonomous Transaction Processing connection to the source Oracle deployment.
7. Assign the Amazon S3 connection to the target Big Data deployment.

Task 2: Create the Extract

1. On the Deployments page, select the source Autonomous Transaction Processing deployment.
2. On the deployment details page, click **Launch Console**.
3. Log in with the source deployment's administrator username and password.
4. Add transaction information.
5. Add an Extract.

Task 3: Create the Distribution Path

1. Create a user for the Distribution Path in the target Big Data deployment.
 - a. On the Deployments page, select the target deployment to view its details.
 - b. On the deployment details page, click **Launch Console**.
 - c. Sign in to the target Big Data deployment console using the Administrator username and password.
 - d. In the Big Data deployment console, open the navigation menu, and then click **Administrator**.
 - e. On the Users page, click **Add New User** (plus icon).
 - f. Complete the fields as follows, and then click **Submit**.
 - i. For **Username**, enter a name, such as `ggsnet`.
 - ii. From the **Role** dropdown, select **Operator**.

- iii. Enter a password twice for verification.

The new user appears in the Users list.

2. In the source ATP GoldenGate deployment console, add a credential for the user created in Step 1.
 - a. In the source ATP GoldenGate deployment console, click Administration Service, open the navigation menu, and then select **Configuration**.
 - b. On the Credentials page, click **Add Credential**, and then complete the fields as follows:
 - i. For **Credential Domain**, enter `GGSNetwork`.
 - ii. For **Credential Alias**, enter `dpuser`.
 - iii. For **User ID**, enter the name of the user created in Step 1 (`ggsnet`).
 - iv. Enter the user's password twice for verification.
 - c. Click **Submit**.

The credential appears in the Credentials list.
3. Click **Distribution Service**, and then click **Add Path** (plus icon).
4. Complete the Add Path form fields as follows, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.
 - b. For **Source Extract**, select the Extract created in Task 2.
 - c. For **Source Trail Name**, select the Extract Trail from Task 2.
 - d. For **Target Authentication Method**, select **UserID Alias**.
 - e. For **Target protocol**, select `wss`.
 - f. For **Target Host**, enter the host domain of the target deployment.

 **Note:**

You can copy and paste the URL of the target Big Data deployment console and remove the `https://` and any trailing slashes.

- g. For **Port Number**, enter `443`.
 - h. For **Trail Name**, enter a two-character name, such as `E1`.
 - i. For **Domain**, enter the name of the Credential Domain created in Step 2 (`GGSNetwork`).
 - j. For **Alias**, enter the Credential Alias created in Step 2 (`dpuser`).
- You return to the Distribution Service Overview page where you can view the status of the created path.
5. In the target Big Data deployment console, review the Receiver Path.
 - a. In the target Big Data deployment console, click **Receiver Service**.
 - b. Review the path details. This path was created as a result of the Distribution Path created in the previous step.

Task 4: Add the Replicat

To add a Replicat for an Amazon S3 :

1. In the target Big Data deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, for Replicat type, select either **Classic** or **Coordinated**, and then click **Next**.
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a short description to distinguish this process from others.
 - c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **Amazon S3** from the dropdown.
 - e. For **Available aliases** for Amazon S3, select your alias from the dropdown.
4. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP SRC_OCIGLL.*, TARGET *.*;
```

5. On the Properties File page, configure the S3 Event Handler properties as needed. Some properties to consider modifying include:
 - a. `gg.eventhandler.s3.region`: provide the AWS region for the target S3 bucket
 - b. `gg.eventhandler.s3.bucketMappingTemplate`: provide target S3 bucket name. If bucket does not exist, it can be auto created by OCI GoldenGate. You can provide static bucket names or you can use Template Keywords to assign bucket names dynamically.
 - c. (Optional) `gg.eventhandler.s3.format`: Select how to format the output. JSON is the default setting. Available options include:
 - i. `delimitedtext`
 - ii. `json`
 - iii. `json_row`
 - iv. `xml`
 - v. `avro_row_ocf`
 - vi. `avro_op_ocf`
6. Click **Create and Run**.

You return to the Overview page, where you can review the Replicat details.

Task 5: Verify the replication

Perform some updates to the source Autonomous Transaction Processing instance to verify replication to Amazon S3.

1. In the Oracle Cloud console, open the navigation menu, select **Oracle Database**, and then select **Autonomous Transaction Processing**.
2. In the list of ATP instances, select your source instance to view its details.

3. On the database details page, click **Database actions**.
4. On the Database actions home page, select **SQL**.
5. Enter the following into the worksheet and click **Run Script**:

```
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1001,0,'Brendt','Paul','10 Jasper Blvd.','107','(212)
555 2146',19,10);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1002,0,'McCarthy','Robin','27 Pasadena
Drive',11,'(214) 555 3075',29,11);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1003,0,'Travis','Peter','7835 Hartford
Drive',12,'(510) 555 4448',34,12);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1004,0,'Larson','Joe','87 Carmel Blvd.','13','(213)
555 5095',45,13);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1005,0,'Goldschmidt','Tony','91 Torre
drive',14,'(619) 555 6529',55,20);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1006,0,'Baker','William','2890 Grant
Avenue',15,'(312) 555 7040',64,21);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1007,0,'Swenson','Jack','64 Imagination
Drive',19,'(202) 555 8125',74,22);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1008,0,'Brendt','Paul','10 Jasper Blvd.','107','(212)
555 2146',19,10);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1009,0,'McCarthy','Robin','27 Pasadena
Drive',11,'(214) 555 3075',29,11);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1010,0,'Travis','Peter','7835 Hartford
Drive',12,'(510) 555 4448',34,12);
```

6. In the source ATP deployment console, select the Extract name, and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CUSTOMER has 10 inserts.
7. In the target Big Data deployment console, select the Replicat name, and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CUSTOMER has 10 inserts

Task 6: Monitor and maintain processes

1. Monitor performance.

2. Manage Trail files.

Replicate data from ATP to Kafka

Learn to configure OCI GoldenGate to replicate data from Autonomous Transaction Processing to Apache Kafka.

Overview

Using OCI GoldenGate, you can replicate data from an Autonomous Database to a Big Data target, like Apache Kafka. This quickstart demonstrates how to set up an OCI GoldenGate deployment, source and target connections, and replication processes.

Before you begin

To successfully complete this quickstart, you must have the following:

- A VCN with port 9092 (or another port used for the bootstrap server) open in the Ingress rules.
- A source Autonomous Transaction Processing instance
- A target Apache Kafka node

Task 0: Set up the source Autonomous Database

If you don't already have a source database set up for replication, you can follow these steps to load a sample schema to use for this quickstart. This quickstart uses Autonomous Transaction Processing (ATP) for the source database.

To set up the source Autonomous Database:

1. Download and unzip the [sample database schema](#).
2. In the Oracle Cloud console, select your ATP instance from the Autonomous Databases page to view its details and access Database Actions.
3. Click **Open Database Actions**.
4. Enable the GGADMIN user:
 - a. Under **Administration**, click **Database Users**.
 - b. Locate **GGADMIN** and then click its ellipsis menu (three dots) and select **Edit**.
 - c. In the Edit User panel, enter the GGADMIN password, confirm the password, and then disable **Account is Locked**.
 - d. Click **Apply Changes**.
5. Load the source sample schema and data:
 - a. From the Database Actions Selector menu, under **Development**, select **SQL**.
 - b. Copy and paste the script from **OCIGLL_OCIGGS_SETUP_USERS_ATP.sql** into the SQL worksheet.
 - c. Click **Run Script**. The Script Output tab displays confirmation messages.
 - d. Clear the SQL worksheet and then copy and paste the SQL script from **OCIGLL_OCIGGS_SRC_USER_SEED_DATA.sql**.

**Tip:**

You may need to run each statement separately for the SQL tool to execute the scripts successfully.

- e. To verify that the tables were created successfully, close the SQL window and reopen it again. In the Navigator tab, look for the SRC_OCIGLL schema and then select tables from their respective dropdowns.
6. Enable supplemental logging:
 - a. Clear the SQL Worksheet.
 - b. Enter the following statement, and then click **Run Statement**:

```
ALTER PLUGGABLE DATABASE ADD SUPPLEMENTAL LOG DATA;
```

Task 1: Create the OCI GoldenGate resources

This quickstart example requires deployments and connections for both the source and target.

1. Create an Oracle deployment for the source Autonomous Transaction Processing instance.
2. Create a Big Data deployment for the target Apache Kafka node.
3. Create a connection for the source Autonomous Transaction Processing instance.
4. Create a connection for the target Kafka node.
5. (Optional) If your Big Data deployment doesn't have a public endpoint, then create a connection to GoldenGate, and then assign this connection to the source Oracle deployment.
6. Assign the ATP connection to the source Oracle deployment.
7. Assign the Kafka connection to the target Big Data deployment.

Task 2: Add the Extract

1. On the Deployments page, select the source Autonomous Transaction Processing deployment.
2. On the deployment details page, click **Launch Console**.
3. Log in with the source deployment's administrator username and password.
4. Add transaction information.
5. Add an Extract

Task 3: Add a Distribution Path

1. Create a user for the Distribution Path in the target Big Data deployment console.
 - a. On the Deployments page, select the target deployment to view its details.
 - b. On the deployment details page, click **Launch Console**.

- c. Sign in to the target Big Data deployment console using the Administrator username and password.
- d. In the OCI GoldenGate deployment console, open the navigation menu, and then click **Administrator**.
- e. On the Users page, click **Add New User** (plus icon).
- f. Complete the fields as follows, and then click **Submit**.
 - For **Username**, enter a name, such as `ggsnet`.
 - From the **Role** dropdown, select **Operator**.
 - Enter a password twice for verification.

The new user appears in the Users list.

2. In the source ATP deployment console, add a credential for the user created in Step 1.
 - a. In the source ATP deployment console, click **Administration Service**, open the navigation menu, and then select **Configuration**.
 - b. On the Credentials page, click **Add Credential**, and then complete the fields as follows:
 - For **Credential Domain**, enter `GGSNetwork`.
 - For **Credential Alias**, enter `dpuser`
 - For **User ID**, enter the name of the user created in Step 1 (`ggsnet`)
 - Enter the user's password twice for verification.
 - c. Click **Submit**.

The credential appears in the Credentials list.

3. Click **Distribution Service**, and then click **Add Path** (plus icon).
4. Complete the Add Path form fields as follows, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.
 - b. For **Source Extract**, select the Extract created in Task 2.
 - c. For **Source Trail Name**, select the Extract Trail from Task 2.
 - d. For **Target Authentication Method**, select **UserID Alias**.
 - e. For **Target** protocol, select **wss**.
 - f. For **Target Host**, enter the host domain of the target deployment.

 **Note:**

You can copy and paste the URL of the target OCI GoldenGate deployment console and remove the `https://` and any trailing slashes.

- g. For **Port Number**, enter `443`.
- h. For **Trail Name**, enter a two-character name, such as `E1`.
- i. For **Domain**, enter the name of the Credential Domain created in Step 2 (`GGSNetwork`).
- j. For **Alias**, enter the Credential Alias created in Step 2 (`dpuser`).

You're returned to the Distribution Service Overview page where you can view the status of the created path.

5. In the target Big Data deployment console, review the Receiver Path.
 - a. In the target Big Data OCI GoldenGate deployment console, click **Receiver Service**.
 - b. Review the path details. This path was created as a result of the Distribution Path created in the previous step.

Task 4: Add the Replicat

1. In the target Big Data deployment console, click **Administrator Service**, and then click **Add Replicat** (plus icon).
2. On the Add Replicat page, under Replicat type, select **Classic Replicat**, and then click **Next**.
3. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name.
 - For **Trail Name**, enter the name of the Trail from Task 2.
 - For **Target**, select **Kafka**.
 - For **Alias**, select the Kafka connection created in Task 1.
4. On the Replicat Parameters page, leave the default, and then click **Next**:
5. On the Properties page, replace `<stream-name>` with the name of your Stream for the following property:

```
gg.handler.oss.topicMappingTemplate=<stream-name>
```
6. Click **Create and Run**.

You're returned to the Overview page, where you can review the Replicat details.

Task 5: Verify the replication

Perform some updates to the source Autonomous Transaction Processing instance to verify replication to Apache Kafka.

1. In the Oracle Cloud console, open the navigation menu, select **Oracle Database**, and then select **Autonomous Transaction Processing**.
2. In the list of Autonomous Transaction Processing instances, select your source instance to view its details.
3. On the database details page, click **Database Actions**.

Note:

You should be automatically logged in. If not, log in with the database credentials.

4. On the Database Actions home page, select **SQL**.

5. Enter the following into the worksheet and click **Run Script**.

```
Insert into SRC_OCIGLLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1000, 'Houston', 20, 743113);
Insert into SRC_OCIGLLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1001, 'Dallas', 20, 822416);
Insert into SRC_OCIGLLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1002, 'San Francisco', 21, 157574);
Insert into SRC_OCIGLLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1003, 'Los Angeles', 21, 743878);
Insert into SRC_OCIGLLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1004, 'San Diego', 21, 840689);
Insert into SRC_OCIGLLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1005, 'Chicago', 23, 616472);
Insert into SRC_OCIGLLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1006, 'Memphis', 23, 580075);
Insert into SRC_OCIGLLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1007, 'New York City', 22, 124434);
Insert into SRC_OCIGLLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1008, 'Boston', 22, 275581);
Insert into SRC_OCIGLLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1009, 'Washington D.C.', 22, 688002);
```

6. In the source ATP deployment console, select the Extract name, and then click **Statistics**. Verify that **SRC_OCIGLLL.SRC_CITY** has 10 inserts.
7. In the target Big Data deployment console, select the Replicat name, and then click **Statistics**. Verify that **SRC_OCIGLLL.SRC_CITY** has 10 inserts.

Task 6: Monitor and maintain processes

1. Monitor performance.
2. Manage Trail files.

Replicate data from Autonomous Transaction Processing to Azure Data Lake Storage

This quickstart demonstrates how to replicate data from Autonomous Transaction Processing to Azure Data Lake Storage Gen 2 using OCI GoldenGate.

Task 0: Set up the source Autonomous Database

If you don't already have a source database set up for replication, you can follow these steps to load a sample schema to use for this quickstart. This quickstart uses Autonomous Transaction Processing (ATP) for the source database.

To set up the source Autonomous Database:

1. Download and unzip the [sample database schema](#).
2. In the Oracle Cloud console, select your ATP instance from the Autonomous Databases page to view its details and access Database Actions.
3. Click **Database actions**.

4. Enable the GGADMIN user:
 - a. Under Administration, click **Database Users**.
 - b. Locate **GGADMIN** and then click its ellipsis menu (three dots) and select **Edit**.
 - c. In the Edit User panel, enter the GGADMIN password, confirm the password, and then disable **Account is Locked**.
 - d. Click **Apply Changes**.
5. Load the load the source sample schema and data:
 - a. From the Database Actions Selector menu, under Development, select **SQL**.
 - b. Copy and paste the script from `OCIGLL_OCIGGS_SETUP_USERS_ATP.sql` into the SQL worksheet.
 - c. Click **Run Script**. The Script Output tab displays confirmation messages.
 - d. Clear the SQL worksheet and then copy and paste the SQL script from `OCIGLL_OCIGGS_SRC_USER_SEED_DATA.sql`.

 **Tip:**

You may need to run each statement separately for the SQL tool to execute the scripts successfully.

- e. To verify that the tables were created successfully, close the SQL window and reopen it again. In the Navigator tab, look for the `SRC_OCIGLL` schema and then select tables from their respective dropdowns.
6. Enable supplemental logging:
 - a. Clear the SQL Worksheet.
 - b. Enter the following statement, and then click **Run Statement**:

```
ALTER PLUGGABLE DATABASE ADD SUPPLEMENTAL LOG DATA;
```

Task 1: Create the OCI GoldenGate resources

This quickstart example requires deployments and connections for both the source and target.

1. Create an Oracle deployment for the source Autonomous Transaction Processing instance.
2. Create a Big Data deployment for the target Azure Data Lake Storage.
3. Create a connection for the source Autonomous Transaction Processing instance.
4. Create connection for Azure Data Lake Storage.
5. (Optional) If your Big Data deployment doesn't have a public endpoint, then create a connection to GoldenGate, and then assign this connection to the source Oracle deployment.
6. Assign the Autonomous Transaction Processing connection to the source Oracle deployment.
7. Assign Azure Data Lake Storage connection to the target Big Data deployment.

Task 2: Add and run the Extract

1. On the Deployments page, select the source Autonomous Transaction Processing deployment.
2. On the deployment details page, click **Launch Console**.
3. Log in with the source deployment's administrator username and password.
4. Add transaction information.
5. Add an Extract.

Task 3: Add and run a Distribution Path

1. Create a user for the Distribution Path in the target Big Data deployment console.
 - a. On the Deployments page, select the target deployment to view its details.
 - b. On the deployment details page, click **Launch Console**.
 - c. Sign in to the target Big Data deployment console using the Administrator username and password.
 - d. In the OCI GoldenGate deployment console, open the navigation menu, and then click **Administrator**.
 - e. On the Users page, click **Add New User** (plus icon).
 - f. Complete the fields as follows, and then click **Submit**.
 - For **Username**, enter a name, such as `ggsnet`.
 - From the **Role** dropdown, select **Operator**.
 - Enter a password twice for verification.

The new user appears in the Users list.

2. In the source ATP deployment console, add a credential for the user created in Step 1.
 - a. In the source ATP deployment console, click **Administration Service**, open the navigation menu, and then select **Configuration**.
 - b. On the Credentials page, click **Add Credential**, and then complete the fields as follows:
 - For **Credential Domain**, enter `GGSNetwork`.
 - For **Credential Alias**, enter `dpuser`
 - For **User ID**, enter the name of the user created in Step 1 (`ggsnet`)
 - Enter the user's password twice for verification.
 - c. Click **Submit**.

The credential appears in the Credentials list.

3. Click **Distribution Service**, and then click **Add Path** (plus icon).
4. Complete the Add Path form fields as follows, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.

- b. For **Source Extract**, select the Extract created in Task 2.
- c. For **Source Trail Name**, select the Extract Trail from Task 2.
- d. For **Target Authentication Method**, select **UserID Alias**.
- e. For **Target** protocol, select **wss**.
- f. For **Target Host**, enter the host domain of the target deployment.

 **Note:**

You can copy and paste the URL of the target OCI GoldenGate deployment console and remove the `https://` and any trailing slashes.

- g. For **Port Number**, enter 443.
- h. For **Trail Name**, enter a two-character name, such as `E1`.
- i. For **Domain**, enter the name of the Credential Domain created in Step 2 (`GGNetwork`).
- j. For **Alias**, enter the Credential Alias created in Step 2 (`dpuser`).

You're returned to the Distribution Service Overview page where you can view the status of the created path.

- 5. In the target Big Data deployment console, review the Receiver Path.
 - a. In the target Big Data OCI GoldenGate deployment console, click **Receiver Service**.
 - b. Review the path details. This path was created as a result of the Distribution Path created in the previous step.

Task 4: Add and run the Replicat

- 1. In the target Big Data deployment console, click **Administrator Service**, and then click **Add Replicat** (plus icon).
- 2. On the Add Replicat page, under Replicat type, select **Classic Replicat**, and then click **Next**.
- 3. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name.
 - For **Trail Name**, enter the name of the Trail from Task 2.
 - For **Target**, select **Azure Data Lake Storage**.
 - For **Alias**, select the Azure Data Lake Storage connection created in Task 1.
- 4. On the Replicat Parameters page, leave the default, and then click **Next**:

```
MAP SRC_OCIGLL.*, TARGET *.*;
```

- 5. On the Properties page, configure Azure Data Lake Storage properties.
Required Properties:

- `gg.eventhandler.abs.bucketMappingTemplate`: Name of the Azure Data Lake Storage Container. If container is pre-configured, a static container name can be provided. If Azure authentication method permissions are provided, Template Keywords can be used for auto container creation by OCI GoldenGate.

(Optional) Additional properties you may consider adding:

- `gg.handler.abs.format`: Select how to format the output. **JSON** is the default setting. Available options include:
 - `delimitedtext`
 - `json`
 - `json_row`
 - `xml`
 - `avro_row_ocf`
 - `avro_op_ocf`

6. Click **Create and Run**.

You're returned to the Overview page, where you can review the Replicat details.

Task 5: Verify the replication

Perform some updates to the source Autonomous Transaction Processing instance to verify replication to Azure Data Lake Storage.

1. In the Oracle Cloud console, open the navigation menu, select **Oracle Database**, and then select **Autonomous Transaction Processing**.
2. In the list of Autonomous Transaction Processing instances, select your source instance to view its details.
3. On the database details page, click **Database Actions**.

 **Note:**

You should be automatically logged in. If not, log in with the database credentials.

4. On the Database Actions home page, select **SQL**.
5. Enter the following into the worksheet and click **Run Script**.

```
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
 values (1001,0,'Brendt','Paul','10 Jasper Blvd.',107,'(212) 555
2146',19,10);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
 values (1002,0,'McCarthy','Robin','27 Pasadena Drive',11,'(214) 555
3075',29,11);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
 values (1003,0,'Travis','Peter','7835 Hartford Drive',12,'(510) 555
4448',34,12);
```

```
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1004,0,'Larson','Joe','87 Carmel Blvd.',13,'(213)
555 5095',45,13);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1005,0,'Goldschmidt','Tony','91 Torre drive',14,
'(619) 555 6529',55,20);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1006,0,'Baker','William','2890 Grant Avenue',15,
'(312) 555 7040',64,21);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1007,0,'Swenson','Jack','64 Imagination Drive',19,
'(202) 555 8125',74,22);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1008,0,'Brendt','Paul','10 Jasper Blvd.',107,'(212)
555 2146',19,10);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1009,0,'McCarthy','Robin','27 Pasadena Drive',11,
'(214) 555 3075',29,11);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1010,0,'Travis','Peter','7835 Hartford Drive',12,
'(510) 555 4448',34,12);
```

6. In the source GoldenGate OCI GoldenGate deployment console, select the Extract name, and then click **Statistics**. Verify that **SRC_OCIGLL.SRC_CUSTOMER** has 10 inserts.
7. In the target Big Data OCI GoldenGate deployment console, select the Replicat name, and then click **Statistics**. Verify that **SRC_OCIGLL.SRC_CUSTOMER** has 10 inserts.
8. In Azure console, navigate to Azure BLOB Storage container and check its contents.

Task 6: Monitor and maintain processes

1. Monitor performance.
2. Manage Trail files.

Replicate data from Autonomous Transaction Process to Azure Synapse

This quickstart demonstrates how to set up a data replication from Autonomous Transaction Processing to Azure Synapse Database using OCI GoldenGate.

Before you begin

To successfully complete this quickstart, you must have the following:

- [Azure Synapse Workspace](#)
- Azure Storage Container associated to Azure Synapse Workspace
- Azure Synapse [Database Scoped Credential](#) to give permissions to SQL pool to access Storage Account
- Target table existence: The target tables should exist on the Synapse database before replication. As GoldenGate uses Merge SQL Statement, the target table must be a hash distributed table.
- OCI GoldenGate Azure Data Lake Storage Connection assigned to deployment. If it doesn't exist, create an Azure Data Lake Storage connection and assign to deployment.

Task 0: Set up the source Autonomous Database

If you don't already have a source database set up for replication, you can follow these steps to load a sample schema to use for this quickstart. This quickstart uses Autonomous Transaction Processing (ATP) for the source database.

To set up the source Autonomous Database:

1. Download and unzip the [sample database schema](#).
2. In the Oracle Cloud console, select your ATP instance from the Autonomous Databases page to view its details and access Database Actions.
3. Click Open **Database Actions**.
4. Enable the GGADMIN user:
 - a. Under Administration, click **Database Users**.
 - b. Locate **GGADMIN** and then click its ellipsis menu (three dots) and select **Edit**.
 - c. In the Edit User panel, enter the GGADMIN password, confirm the password, and then disable **Account is Locked**.
 - d. Click **Apply Changes**.
5. Load the load the source sample schema and data:
 - a. From the Database Actions Selector menu, under Development, select **SQL**.
 - b. Copy and paste the script from `OCIGLL_OCIGGS_SETUP_USERS_ATP.sql` into the SQL worksheet.
 - c. Click Run Script. The Script Output tab displays confirmation messages.
 - d. Clear the SQL worksheet and then copy and paste the SQL script from `OCIGLL_OCIGGS_SRC_USER_SEED_DATA.sql`.

Tip:

You may need to run each statement separately for the SQL tool to execute the scripts successfully.

- e. To verify that the tables were created successfully, close the SQL window and reopen it again. In the Navigator tab, look for the `SRC_OCIGLL` schema and then select tables from their respective dropdowns.
6. Enable supplemental logging.

- a. Clear the SQL Worksheet.
- b. Enter the following statement, and then click **Run Statement**:

```
ALTER PLUGGABLE DATABASE ADD SUPPLEMENTAL LOG DATA;
```

Task 1: Create the OCI GoldenGate resources

This quickstart example requires deployments and connections for both the source and target.

1. Create an Oracle deployment for the source Autonomous Transaction Processing instance.
2. Create a Big Data deployment for the target Azure Synapse.
3. Create a connection to the source Autonomous Transaction Processing.
4. Create an Azure Synapse Analytics connection.
5. (Optional) If your Big Data deployment doesn't have a public endpoint, then create a connection to GoldenGate, and then assign this connection to the source Oracle deployment.
6. Assign the Autonomous Transaction Processing connection to the source Oracle deployment.
7. Assign Azure Synapse connection to the target Big Data deployment.

Task 2: Add and run the Extract

1. On the Deployments page, select the source Autonomous Transaction Processing deployment.
2. On the deployment details page, click **Launch Console**.
3. Log in with the source deployment's administrator username and password.
4. Add transaction information.
5. Add an Extract.

Task 3: Add and run a Distribution Path

1. Create a user for the Distribution Path in the target Big Data deployment.
 - a. On the Deployments page, select the target deployment to view its details.
 - b. On the deployment details page, click **Launch Console**.
 - c. Sign in to the target OCI GoldenGate deployment console using the Administrator username and password.
 - d. In the OCI GoldenGate deployment console, open the navigation menu, and then click **Administrator**.
 - e. On the Users page, click **Add New User** (plus icon).
 - f. Complete the fields as follows, and then click **Submit**
 - For **Username**, enter a name, such as `ggsnet`.

- From the **Role** dropdown, select **Operator**.
- Enter a password twice for verification.

The new user appears in the Users list.

2. In the source ATP GoldenGate deployment console, add a credential for the user created in Step 1.
 - a. In the source ATP GoldenGate deployment console, click **Administration Service**, open the navigation menu, and then select **Configuration**.
 - b. On the Credentials page, click **Add Credential**, and then complete the fields as follows
 - For **Credential Domain**, enter `GGSNetwork`.
 - For **Credential Alias**, enter `dpuser`
 - For **User ID**, enter the name of the user created in Step 1 (`ggsnet`)
 - Enter the user's password twice for verification.
 - c. Click **Submit**.

The credential appears in the Credentials list.

3. Click **Distribution Service**, and then click **Add Path** (plus icon).
4. Complete the Add Path form fields as follows, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.
 - b. For **Source Extract**, select the Extract created in Task 2.
 - c. For **Source Trail Name**, select the Extract Trail from Task 2.
 - d. For **Target Authentication Method**, select **UserID Alias**.
 - e. For **Target protocol**, select **wss**.
 - f. For **Target Host**, enter the host domain of the target deployment.

 **Note:**

You can copy and paste the URL of the target OCI GoldenGate deployment console and remove the `https://` and any trailing slashes.

- g. For **Port Number**, enter `443`.
- h. For **Trail Name**, enter a two-character name, such as `E1`.
- i. For **Domain**, enter the name of the Credential Domain created in Step 2 (`GGSNetwork`).
- j. For **Alias**, enter the Credential Alias created in Step 2 (`dpuser`).

You're returned to the Distribution Service Overview page where you can view the status of the created path.

5. In the target Big Data deployment console, review the Receiver Path.
 - a. In the target Big Data deployment console, click **Receiver Service**.
 - b. Review the path details. This path was created as a result of the Distribution Path created in the previous step.

Task 4: Add and run the Replicat

1. In the target Big Data deployment console, click **Administrator Service**, and then click **Add Replicat** (plus icon).
2. On the Add Replicat page, under Replicat type, select **Classic Replicat**, and then click **Next**.
3. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name.
 - For **Trail Name**, enter the name of the Trail from Task 2.
 - For **Target**, select **Azure Synapse Analytics**.
 - For **Alias**, select the Azure Data Lake Storage connection created in Task 1.
4. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP SRC_OCIGLL.SRC_CUSTOMER, TARGET dbo.SRC_CUSTOMER;
```

5. On the Properties page, configure Azure Synapse properties:

 **Note:**

Edit the properties marked as `TODO`.

- `gg.eventhandler.abs.bucketMappingTemplate` is the Azure Storage Container associated with the Azure Synapse Workspace
 - `gg.eventhandler.synapse.credential` is the name of the credential used to authenticate Azure Storage Container associated with the Azure Synapse Workspace
6. Click **Create and Run**.

You're returned to the Overview page, where you can review the Replicat details.

Task 5: Verify the replication

Perform some updates to the source Autonomous Transaction Processing instance to verify replication to Azure Data Lake Storage.

1. In the Oracle Cloud console, open the navigation menu, select **Oracle Database**, and then select **Autonomous Transaction Processing**.
2. In the list of Autonomous Transaction Processing instances, select your source instance to view its details.
3. On the database details page, click **Database Actions**.

 **Note:**

You should be automatically logged in. If not, log in with the database credentials.

4. On the Database Actions home page, select **SQL**.
5. Enter the following into the worksheet and click **Run Script**.
6. In the source GoldenGate OCI GoldenGate deployment console, select the Extract name, and then click **Statistics**. Verify that **SRC_OCIGLL.SRC_CUSTOMER** has 7 inserts.

```

Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
values (1001,0,'Brendt','Paul','10 Jasper Blvd.','107','(212) 555
2146',19,10);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
values (1002,0,'McCarthy','Robin','27 Pasadena Drive',11,'(214) 555
3075',29,11);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
values (1003,0,'Travis','Peter','7835 Hartford Drive',12,'(510) 555
4448',34,12);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
values (1004,0,'Larson','Joe','87 Carmel Blvd.','13','(213) 555
5095',45,13);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
values (1005,0,'Goldschmidt','Tony','91 Torre drive',14,'(619) 555
6529',55,20);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
values (1006,0,'Baker','William','2890 Grant Avenue',15,'(312) 555
7040',64,21);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
values (1007,0,'Swenson','Jack','64 Imagination Drive',19,'(202) 555
8125',74,22)

```

7. In the target Big Data OCI GoldenGate deployment console, select the Replicat name, and then click **Statistics**. Verify that **SRC_OCIGLL.SRC_CUSTOMER** has 7 inserts.
8. In Azure console, navigate to Azure Synapse workspace Console. Run `Select *` from **dbo.SRC_CUSTOMER** and verify that **SRC_OCIGLL.SRC_CUSTOMER** has 7 inserts.

Replicate data from Autonomous Transaction Processing to Confluent Kafka

This quickstart demonstrates how to replicate data from Autonomous Transaction Processing to Confluent Kafka using OCI GoldenGate.

Before you begin

To successfully complete this quickstart, you must have the following:

- A source Autonomous Transaction Processing instance.
- A Confluent Platform instance.

Task 0.1: Set up the source Autonomous Database

If you don't already have a source database set up for replication, you can follow these steps to load a sample schema to use for this quickstart. This quickstart uses Autonomous Transaction Processing (ATP) for the source database.

To set up the source Autonomous Database:

1. Download and unzip the [sample database schema](#).
2. In the Oracle Cloud console, select your ATP instance from the Autonomous Databases page to view its details and access Database Actions.
3. Click Open **Database Actions**.
4. Enable the GGADMIN user:
 - a. Under Administration, click **Database Users**.
 - b. Locate **GGADMIN** and then click its ellipsis menu (three dots) and select Edit.
 - c. In the Edit User panel, enter the GGADMIN password, confirm the password, and then disable **Account is Locked**.
 - d. Click **Apply Changes**.
5. Load the load the source sample schema and data:
 - a. From the Database Actions Selector menu, under Development, select **SQL**.
 - b. Copy and paste the script from `OCIGLL_OCIGGS_SETUP_USERS_ATP.sql` into the SQL worksheet.
 - c. Click **Run Script**. The Script Output tab displays confirmation messages.
 - d. Clear the SQL worksheet and then copy and paste the SQL script from `OCIGLL_OCIGGS_SRC_USER_SEED_DATA.sql`.

 **Tip:**

You may need to run each statement separately for the SQL tool to execute the scripts successfully.

- e. To verify that the tables were created successfully, close the SQL window and reopen it again. In the Navigator tab, look for the `SRC_OCIGLL` schema and then select tables from their respective dropdowns.
6. Enable supplemental logging:
 - a. Clear the SQL Worksheet.
 - b. Enter the following statement, and then click **Run Statement**:

```
ALTER PLUGGABLE DATABASE ADD SUPPLEMENTAL LOG DATA;
```

Task 0.2: Set up the target Confluent Cloud

If you don't already have a target Confluent Cloud cluster, topic and schema, you can do the following:

1. [Create a cluster](#).
2. [Create an API Key for the Cluster](#). Note the API Key and Secret for the next steps.
3. [Enable Schema Registry](#), and then [create an API Key for Confluent Cloud Schema Registry](#). Note the API Key and Secret for the next steps.
4. [Create a topic](#) in the cluster. Note the name of the topic for the next steps.

Task 1: Create the OCI GoldenGate resources

This quickstart example requires deployments and connections for both the source and target.

1. Create an Oracle deployment for the source Autonomous Transaction Processing instance.
2. Create a Big Data deployment for the target Confluent Cloud.
3. Create a connection to the source Autonomous Transaction Processing.
4. Create a Confluent Kafka connection.
5. Create a Confluent Schema Registry.
6. (Optional) If your Big Data deployment doesn't have a public endpoint, then create a connection to GoldenGate, and then assign this connection to the source Oracle deployment.
7. Assign the Autonomous Transaction Processing connection to the source Oracle deployment.
8. Assign Confluent Kafka and Confluent Schema Registry connections to the target Big Data deployment.

Task 2: Add the Extract

1. On the Deployments page, select the source Autonomous Transaction Processing deployment.
2. On the deployment details page, click **Launch Console**.
3. Log in with the source deployment's administrator username and password.
4. Add transaction information.
5. Add an Extract.

Task 3: Add a Distribution Path

1. Create a user for the Distribution Path in the target Big Data deployment.
 - a. On the Deployments page, select the target deployment to view its details.
 - b. On the deployment details page, click **Launch Console**.

- c. Sign in to the target OCI GoldenGate deployment console using the Administrator username and password.
- d. In the OCI GoldenGate deployment console, open the navigation menu, and then click **Administrator**.
- e. On the Users page, click **Add New User** (plus icon).
- f. Complete the fields as follows, and then click **Submit**.
 - For **Username**, enter a name, such as `ggsnet`.
 - From the **Role** dropdown, select **Operator**.
 - Enter a password twice for verification.

The new user appears in the Users list.

2. In the source ATP GoldenGate deployment console, add a credential for the user created in Step 1.
 - a. In the source ATP GoldenGate deployment console, click **Administration Service**, open the navigation menu, and then select **Configuration**.
 - b. On the Credentials page, click **Add Credential**, and then complete the fields as follows:
 - For **Credential Domain**, enter `GGSSNetwork`.
 - For **Credential Alias**, enter `dpuser`
 - For **User ID**, enter the name of the user created in Step 1 (`ggsnet`)
 - Enter the user's password twice for verification.

- c. Click **Submit**.

The credential appears in the Credentials list.

3. Click **Distribution Service**, and then click **Add Path** (plus icon).
4. Complete the Add Path form fields as follows, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.
 - b. For **Source Extract**, select the Extract created in Task 2.
 - c. For **Source Trail Name**, select the **Extract Trail** from Task 2.
 - d. For **Target Authentication Method**, select **UserID Alias**.
 - e. For **Target protocol**, select `wss`.
 - f. For **Target Host**, enter the host domain of the target deployment.

 **Note:**

You can copy and paste the URL of the target OCI GoldenGate deployment console and remove the `https://` and any trailing slashes.

- g. For **Port Number**, enter `443`.
- h. For **Trail Name**, enter a two-character name, such as `E1`.
- i. For **Domain**, enter the name of the Credential Domain created in Step 2 (`GGSSNetwork`).

- j. For **Alias**, enter the Credential Alias created in Step 2 (dpuser).

You're returned to the Distribution Service Overview page where you can view the status of the created path.

5. In the target Big Data deployment console, review the **Receiver Path**.
 - a. In the target Big Data OCI GoldenGate deployment console, click **Receiver Service**.
 - b. Review the path details. This path was created as a result of the Distribution Path created in the previous step.

Task 4: Add and run the Replicat

1. In the target Big Data deployment console, click **Administrator Service**, and then click **Add Replicat** (plus icon).
2. On the Add Replicat page, under Replicat type, select **Classic Replicat**, and then click **Next**.
3. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name.
 - For **Trail Name**, enter the name of the Trail from Task 2.
 - For **Target**, select **Kafka**.
 - For **Available Alias**, select the Confluent Kafka connection created in Task 1.
 - Enable **Kafka Connect**.
 - For **Converter**, select **Avro**.
 - For **Schema Registry**, select the Confluent Schema Registry connection created in Task 1.
4. On the Replicat Parameters page, leave the default, and then click **Next**:

```
MAP SRC_OCIGLL.SRC_CUSTOMER, Table SRC.CUSTOMER;
```

5. On the Properties page, configure Kafka Connect properties.

Note:

Edit only the properties marked as `TODO`.

- `gg.handler.kafkaconnect.topicMappingTemplate` is the name of the topic that was created in Task 0.2, Step 4
6. Click **Create and Run**.

You're returned to the Overview page, where you can review the Replicat details.

Task 5: Verify the replication

Perform some updates to the source Autonomous Transaction Processing instance to verify replication to Confluent Kafka Cloud.

1. In the Oracle Cloud console, open the navigation menu, select **Oracle Database**, and then select **Autonomous Transaction Processing**.
2. In the list of Autonomous Transaction Processing instances, select your source instance to view its details.
3. On the database details page, click **Database Actions**.

 **Note:**

You should be automatically logged in. If not, log in with the database credentials.

4. On the Database Actions home page, select **SQL**.
5. Enter the following into the worksheet and click **Run Script**.
6. In the source GoldenGate OCI GoldenGate deployment console, select the Extract name, and then click **Statistics**. Verify that **SRC_OCIGLL.SRC_CUSTOMER** has 7 inserts.

```
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1001,0,'Brendt','Paul','10 Jasper Blvd.','107','(212)
555 2146',19,10);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1002,0,'McCarthy','Robin','27 Pasadena
Drive',11,'(214) 555 3075',29,11);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1003,0,'Travis','Peter','7835 Hartford
Drive',12,'(510) 555 4448',34,12);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1004,0,'Larson','Joe','87 Carmel Blvd.','13','(213)
555 5095',45,13);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1005,0,'Goldschmidt','Tony','91 Torre
drive',14,'(619) 555 6529',55,20);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1006,0,'Baker','William','2890 Grant
Avenue',15,'(312) 555 7040',64,21);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P
RS_ID) values (1007,0,'Swenson','Jack','64 Imagination
Drive',19,'(202) 555 8125',74,22)
```

7. In the target Big Data OCI GoldenGate deployment console, select the Replicat name, and then click **Statistics**. Verify that **SRC_OCIGLL.SRC_CUSTOMER** has 7 inserts.

Task 6: (Optional) Verify the replication in Confluent CLI

1. Install [Confluent CLI](#) in your system.

2. Execute:

```
export PATH=$(pwd)/bin:$PATH
confluent login
```

3. List environments and select your environment

```
confluent environment list
confluent environment use <your_environment_ID>
```

4. List clusters and select your cluster

```
confluent kafka cluster list
confluent kafka cluster use <your_cluster_id>
```

5. Store API Key and Secret locally

```
confluent api-key store <cluster_api_key> <cluster_api_secret>
confluent api-key use <cluster_api_key> --resource <cluster_id>
```

6. List topics

```
confluent kafka topic list
```

7. View messages

```
confluent kafka topic consume --value-format avro --from-beginning
<topic_name>
```

Replicate data from Amazon RDS to OCI Object Storage

Learn to replicate data from Amazon RDS for Oracle to OCI Object Storage using Oracle Cloud Infrastructure GoldenGate.

Before you begin

To successfully complete this quickstart, you must have the following:

- Sign up or Sign in to your Oracle Cloud account.
- A source Amazon RDS for Oracle instance, explained in Task 0.

Task 0: Set up the source Amazon RDS for Oracle

Provision an Amazon RDS for Oracle instance and carry out following steps to setup a source database for use with Oracle GoldenGate.

1. Turn on supplemental logging on the source database.

2. Set `ENABLE_GOLDENGATE_REPLICATION` initialization parameter to `true`.
3. Enable archiving on Source database and retain archived redo logs.
4. Create an Oracle GoldenGate user account on the source database.
5. Grant user account privileges on the source database.

```
GRANT CREATE SESSION, ALTER SESSION TO GGADMIN;  
GRANT RESOURCE TO GGADMIN;  
GRANT SELECT ANY DICTIONARY TO GGADMIN;  
GRANT FLASHBACK ANY TABLE TO GGADMIN;  
GRANT SELECT ANY TABLE TO GGADMIN;  
GRANT EXECUTE ON DBMS_FLASHBACK TO GGADMIN;  
GRANT SELECT ON SYS.V_$DATABASE TO GGADMIN;  
GRANT ALTER ANY TABLE TO GGADMIN;  
EXEC rdsadmin.rdsadmin_dbms_goldengate_auth.grant_admin_privilege (  
  grantee           => 'GGADMIN',  
  privilege_type    => 'capture',  
  grant_select_privileges => true,  
  do_grants         => TRUE);
```

6. Download and unzip the [sample database schema](#).
7. Load the source sample schema and data:
 - a. Connect to Amazon RDS for Oracle instance from SQL Developer as user `SRC_OCIGLL`.
 - b. Copy and paste the script from `OCIGLL_OCIGGS_SETUP_USERS_ATP.sql` into the SQL worksheet.
 - c. Click **Run Script**. The **Script Output** tab displays confirmation messages.
 - d. Clear the SQL worksheet and then copy and paste the SQL script from `OCIGLL_OCIGGS_SRC_USER_SEED_DATA.sql`.

 **Tip:**

You may need to run each statement separately for the SQL tool to execute the scripts successfully.

- e. To verify that the tables were created successfully, close the SQL window and reopen it again. In the Navigator tab, look for the `SRC_OCIGLL` schema and then select tables from their respective drop-down lists.

Task 1: Create the OCI GoldenGate resources

1. Create an Oracle deployment for the source Amazon RDS for Oracle instance.
2. Create a Big Data deployment target OCI Object Storage bucket.
3. Create a connection for the source Amazon RDS for Oracle instance.
4. Create a connection for the target OCI Object Storage.
5. (Optional) If your Big Data deployment does not have a public endpoint, then create a connection to GoldenGate, and assign this connection to the source Oracle deployment.

6. Assign the Amazon RDS for Oracle connection to the source Oracle deployment.
7. Assign the OCI Object Storage connection to the target Big Data deployment.

Task 2: Add and run the Extract

1. On the Deployments page, select the source Amazon RDS for Oracle deployment.
2. On the deployment details page, click **Launch Console**.
3. Log in with the source deployment's administrator username and password.
4. Add transaction information.
5. Add an Extract.

Task 3: Add and run a Distribution Path

1. Create a user for the Distribution Path in the target Big Data deployment.
 - a. On the Deployments page, select the target deployment to view its details.
 - b. On the deployment details page, click **Launch Console**.
 - c. Sign in to the target OCI GoldenGate deployment console using the Administrator username and password.
 - d. In the OCI GoldenGate deployment console, open the navigation menu, and then click **Administrator**.
 - e. On the Users page, click **Add New User** (plus icon).
 - f. Complete the fields as follows, and then click **Submit**.
 - For **Username**, enter a name, such as `ggsnet`.
 - From the **Role** dropdown, select **Operator**.
 - Enter a password twice for verification.

The new user appears in the Users list.

2. In the Amazon RDS for Oracle GoldenGate deployment console, add a credential for the user created in Step 1.
 - a. In the source Amazon RDS for Oracle GoldenGate deployment console, click **Administration Service**, open the navigation menu, and then select **Configuration**.
 - b. On the Credentials page, click **Add Credential**, and then complete the fields as follows:
 - For **Credential Domain**, enter `GGSNetwork`.
 - For **Credential Alias**, enter `dpuser`
 - For **User ID**, enter the name of the user created in Step 1 (`ggsnet`)
 - Enter the user's password twice for verification.
 - c. Click **Submit**.

The credential appears in the Credentials list.

3. Click **Distribution Service**, and then click **Add Path** (plus icon).
4. Complete the Add Path form fields as follows, and then click **Create and Run**:

- a. For **Path Name**, enter a name.
- b. For **Source Extract**, select the Extract created in Task 2.
- c. For **Source Trail Name**, select the Extract Trail from Task 2.
- d. For **Target Authentication Method**, select **UserID Alias**.
- e. For **Target** protocol, select **wss**.
- f. For **Target Host**, enter the host domain of the target deployment.

 **Note:**

You can copy and paste the URL of the target OCI GoldenGate deployment console and remove the `https://` and any trailing slashes.

- g. For **Port Number**, enter 443.
- h. For **Trail Name**, enter a two-character name, such as `E1`.
- i. For **Domain**, enter the name of the Credential Domain created in Step 2 (`GGSSNetwork`).
- j. For **Alias**, enter the Credential Alias created in Step 2 (`dpuser`).

You're returned to the Distribution Service Overview page where you can view the status of the created path.

5. In the target Big Data deployment console, review the Receiver Path.
 - a. In the target Big Data OCI GoldenGate deployment console, click **Receiver Service**.
 - b. Review the path details. This path was created as a result of the Distribution Path created in the previous step.

Task 4: Add and run the Replicat

1. In the target Big Data deployment console, click **Administrator Service**, and then click **Add Replicat** (plus icon).
2. On the Add Replicat page, under Replicat type, select **Classic Replicat**, and then click **Next**.
3. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name.
 - For **Trail Name**, enter the name of the Trail from Task 2.
 - For **Target**, select **Azure Data Lake Storage**.
 - For **Alias**, select the Azure Data Lake Storage connection created in Task 1.
4. On the Replicat Parameters page, leave the default, and then click **Next**:

```
MAP SRC_OCIGLL.*, TARGET *.*;
```

5. On the Properties page, configure Azure Data Lake Storage properties.
Required Properties:

- `gg.eventhandler.abs.bucketMappingTemplate`: Name of the Azure Data Lake Storage Container. If container is pre-configured, a static container name can be provided. If Azure authentication method permissions are provided, Template Keywords can be used for auto container creation by OCI GoldenGate.

(Optional) Additional properties you may consider adding:

- `gg.handler.abs.format`: Select how to format the output. **JSON** is the default setting. Available options include:
 - `delimitedtext`
 - `json`
 - `json_row`
 - `xml`
 - `avro_row_ocf`
 - `avro_op_ocf`

6. Click **Create and Run**.

You're returned to the Overview page, where you can review the Replicat details.

Task 5: Verify the replication

Perform some updates to the source Autonomous Transaction Processing instance to verify replication to Azure Data Lake Storage.

1. Connect to Amazon RDS for Oracle instance from SQL Developer as user `SRC_OCIGLL`.
2. Enter the following into the worksheet and click **Run Script**.

```
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1000,'Houston',20,743113);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1001,'Dallas',20,822416);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1002,'San Francisco',21,157574);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1003,'Los Angeles',21,743878);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1004,'San Diego',21,840689);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1005,'Chicago',23,616472);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1006,'Memphis',23,580075);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1007,'New York City',22,124434);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1008,'Boston',22,275581);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1009,'Washington D.C.',22,688002);
```

3. In the source GoldenGate OCI GoldenGate deployment console, select the Extract name, and then click **Statistics**. Verify that `SRC_OCIGLL.SRC_CITY` has 10 inserts.
4. In the target Big Data OCI GoldenGate deployment console, select the Replicat name, and then click **Statistics**. Verify that `SRC_OCIGLL.SRC_CITY` has 10 inserts.

Task 6: Monitor and maintain processes

1. Monitor performance.
2. Manage Trail files.

Replicate data from MongoDB to Autonomous JSON Database

Learn to replicate data from MongoDB to Autonomous JSON Database using OCI GoldenGate.

Oracle Autonomous JSON Database is a cloud document database service that makes it simple to develop JSON-centric applications. It features NoSQL-style document APIs (Oracle SODA and Oracle Database API for MongoDB), serverless scaling, high performance ACID transactions, comprehensive security, and low pay-per-use pricing. Learn more about [Autonomous JSON Database](#).

The OCI GoldenGate Big Data deployment type supports no down-time migrations from MongoDB to Autonomous JSON Database. OCI GoldenGate supports both Initial Load Extract and Change Data Capture (CDC) extract from MongoDB.

This quickstart details the process to configure OCI GoldenGate for no down-time migrations from MongoDB to Autonomous JSON Database.

Before you begin

To successfully complete this quickstart, you must have the following:

- MongoDB replica set configured.
 - OCI GoldenGate Big Data capture uses operations log (oplog) to read the CDC records. The oplog is a capped collection that keeps a rolling record of all operations that modify the data stored in your databases. Oplog files are created in MongoDB when Replicat set is enabled. MongoDB Atlas comes with a preconfigured Replicat set configuration. For on premises MongoDB, you need to [deploy a replica set](#).
- MongoDB 3.6 & later.
- OCI GoldenGate support for capture of following operations: INSERT, UPDATE, DELETE.

To learn more, see [Using Oracle GoldenGate Capture for MongoDB](#).

Task 1: Create the OCI GoldenGate resources

This quickstart example requires deployments and connections for both the source and target.

1. [Create an Oracle Autonomous JSON Database](#).
2. Create a Big Data deployment for the source MongoDB and target Oracle Autonomous JSON Database.
3. Create a MongoDB connection.
4. Create connection for target Create an Autonomous JSON Database connection.

5. (Optional) If your Big Data deployment doesn't have a public endpoint, then create a connection to GoldenGate, and then assign this connection to the Big Data deployment.
6. Assign MongoDB and Autonomous JSON connections to the Big Data deployment.

Task 2: Create and run the Change Data Capture Extract for MongoDB

1. In the OCI GoldenGate deployment console, ensure that you're on the Administration Service Overview page, and then click **Add Extract** (plus icon).
2. For Source, select **MongoDB** from the dropdown and select the Extract Type as **Change Data Capture Extract**.
3. On the Extract Options page, complete the following fields, and then click **Next**:
 - a. For **Process name**, enter `MCDC`.
 - b. For **Connection Alias**, select the name of your MongoDB connection from the dropdown.
 - c. For **Trail Name**, enter `M1`.
4. On the Parameter File page, ensure the source mapping includes `TABLE source.*;`
5. Click **Create and Run**.

Task 3: Create and run the Initial Load Extract for MongoDB

1. In the OCI GoldenGate deployment console, ensure that you're on the Administration Service Overview page, and then click **Add Extract** (plus icon).
2. For Source, select **MongoDB** from the dropdown and then select **Initial Load Extract** for Extract type.
3. On the Extract Options page, complete the following and then click **Next**:
 - a. For **Process Name**, enter `MIL`.
 - b. For **Connection Alias**, select the name of your MongoDB connection from the dropdown.
 - c. For **Trail Name**, enter `I1`.
4. On the Parameter File page, make the following changes:
 - a. Locate the `EXTTRAIL` line and replace `$extfilePath` with your trail name. For example:

```
EXTFILE I1
```

- b. Define the source mapping as `TABLE source.*;`

Note:

This is the source database/collection mapping. `TABLE *.*` results in extracting from all the databases/collections.

5. Click **Create and Run**.

When MongoDB Initial Load Extract runs successfully, you will see the statistics in the extract report file.

Task 4: Create and run the Initial Load Replicat

1. In the OCI GoldenGate deployment console, ensure that you're on the Administration Service Overview page, and then click **Add Replicat** (plus icon).
2. Add the Initial Load Replicat.
3. On the Add Replicat page, under Replicat type, select either **Classic** or **Coordinated**, and then click **Next**.
4. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - a. For **Process Name**, enter a name, such as `RIL`.
 - b. For **Trail Name**, enter the name of the Trail from Task 3 (`I1`).
 - c. For **Target**, select the target **ORACLE_AUTONOMOUS_JSON_DATABASE** from the dropdown.
 - d. For **Available Aliases**, select **Autonomous JSON connection** from the dropdown.
5. On the Replicat Parameters page, you can specify parameters to further configure your Replicat, and then click **Next**:

```
MAP *.* , TARGET *.*;
```

6. On the Properties page, review the properties, and then click **Create and Run**.

You return to the Overview page, where you can review the Replicat details. When replicat starts successfully, you'll see it in a running state with a green check. You can review the replicat details and statistics to confirm the replication.

Task 5: Create and run the Change Data Capture Replicat

1. In the OCI GoldenGate deployment console, ensure that you're on the Administration Service Overview page, and then click **Add Replicat** (plus icon).
2. On the Add Replicat page, under Replicat type, select either **Classic** or **Coordinated**, and then click **Next**.
3. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - a. For **Process Name**, enter a name, such as `MCDC`.
 - b. For **Trail Name**, enter the name of the Trail from Task 2 (`M1`).
 - c. For **Target**, select **MongoDB** from the dropdown.
 - d. For **Available Aliases**, select assigned MongoDB connection from the dropdown.
4. On the Parameter File page, you specify parameters to further configure your Replicat, and then click **Next**:

```
HANDLECOLLISIONS  
MAP *.* , TARGET *.*;
```


 **Note:**

Add `HANDLECOLLISIONS` to resolve the issues with duplicate or missing records while applying the replicat.

`HANDLECOLLISIONS` parameter has a negative impact on the performance and that's why it is recommended to use as needed. When your source & target is synced, you can stop the CDC replicat, remove `HANDLECOLLISIONS` and re-start the replicat.

5. On the Properties page, review the properties, and then click **Create and Run**.

You return to the Overview page, where you can review the Replicat details. When Replicat starts successfully, you'll see it in a running state with a green check. You can check the Replicat details and statistics confirm the replication.

Capture data from Kafka platforms

Use OCI GoldenGate to extract message from Kafka platform streaming sources.

Overview

You can use OCI GoldenGate to capture messages from the following streaming sources:

- Apache Kafka
- OCI Streaming
- Confluent Kafka, with or without Schema Registry
- Azure Event Hubs
- Amazon MSK

OCI GoldenGate reads messages from a Kafka topic or topics, and then converts the data into logical change records written to GoldenGate Trail files. GoldenGate Replicat processes can then use the generate Trail files to propagate data to support RDBMS implementations.

Task 1: Configure Consumer properties

1. Create a Kafka Consumer properties file with one of the following deserializers or converters. If the source is a topic in Confluent Kafka with Schema Registry, you can use the Avro converter. For other sources, use the JSON converter or deserializer as needed:
 - Kafka Consumer properties for JSON deserializer:

```
key.deserializer=org.apache.kafka.common.serialization.ByteArrayDeseriali
alizer
value.deserializer=org.apache.kafka.common.serialization.ByteArrayDese
rializer
```

- Kafka Consumer properties for JSON converter:

```
key.converter=org.apache.kafka.connect.json.JsonConverter
value.converter=org.apache.kafka.connect.json.JsonConverter
```

- Kafka Consumer properties for Avro converter:

```
key.converter=io.confluent.connect.avro.AvroConverter  
value.converter=io.confluent.connect.avro.AvroConverter
```

2. Save the properties file and note its location.

Task 2: Create OCI GoldenGate resources

This task guides you on how to create new resources if they don't yet exist. Ensure that the Big Data deployment you're using is upgraded to the latest version available.

1. Create an OCI GoldenGate deployment for Big Data.
2. Create a connection.

Note:

When creating any of the following connections, ensure that you, click **Show Advanced Options** and then upload the Consumer properties file.

- For Apache Kafka or Amazon MSK, create a Kafka connection.
 - For Confluent Kafka, create a Confluent Kafka connection.
 - For Confluent Schema Registry, create a Confluent Schema Registry connection.
 - For Azure Event Hubs, create an Azure Event Hubs connection.
 - For OCI Streaming, create an OCI Streaming connection.
3. Assign the connection to the Big Data deployment.

Task 3: Create the Extract

1. Select the Big Data deployment on the Deployments page.
2. On the deployment details page, click **Launch console**.
3. Log in to the Big Data deployment using the credentials specified when you created the deployment in Task 2 Step 1.
4. Add an Extract.
 - a. On the Administration Service Overview page, click **Add Extract** (plus icon).
 - b. On the Add Extract page, for **Extract type**, select **Change Data Capture**, and then click **Next**.
 - c. On the Extract Options page, complete the fields as follows, and then click **Next**:
 - For **Process Name**, enter a name for the extract.
 - For **Alias**, select the connection assigned to the deployment.
 - For **Begin**, select **Now**.
 - For **Trail Name**, enter a 2-character name.

- (Optional) Enable **Kafka Connect**, if the source is a Kafka Connect framework.
- (Optional) Select a **Converter**. If you select **Avro**, select Schema Registry.
- d. On the Parameter File page, leave the table mapping as `TABLE TESTSCHEMA.*`; to listen to all topics in the given bootstrap server. You can also set the table mapping as `TABLE TESTSCHEMA.<topic-name>`; to capture from a designated topic.
- e. Click **Create and Run**.

You return to the Administration Service Overview page, where you can observe the Extract process start and review event messages.

Stage and merge data into Autonomous Data Warehouse using OCI GoldenGate

This quickstart guides you on how to stage and merge data from Autonomous Transaction Processing to Autonomous Data Warehouse using an OCI GoldenGate Big Data deployment.

Before you begin

You must have the following in order to proceed:

- An existing source database.
- An existing target Autonomous Database.
- An existing OCI Object Storage Bucket which will be used as a temporary staging area.
- Before configuring ADW Stage & Merge replication, target schemas and tables should be created in the target ADW instance.
- You can download [Archive.zip](#) and follow Task 0 to set up source and target databases using Autonomous Database.

Task 0: Set up the source & target Autonomous Databases

1. Download and unzip the [sample database schema](#).
2. Set up the source Autonomous Database:
 - a. In the Oracle Cloud console, select your ATP instance from the Autonomous Databases page to view its details and access Database Actions.
 - b. Click **Database Actions**.
 - c. Enable the GGADMIN user:
 - i. Under **Administration**, click Database Users.
 - ii. Locate GGADMIN and then click its ellipsis menu (three dots) and select **Edit**.
 - iii. In the Edit User panel, enter the GGADMIN password, confirm the password, and then disable **Account is Locked**.
 - iv. Click **Apply Changes**.
 - d. Load the load the source sample schema and data:
 - i. From the Database Actions Selector menu, under Development, select **SQL**.
 - ii. Copy and paste the script from `OCIGLL_OCIGGS_SETUP_USERS_ATP.sql` into the SQL worksheet.
 - iii. Click **Run Script**. The Script Output tab displays confirmation messages.

- iv. Clear the SQL worksheet and then copy and paste the SQL script from `OCIGLL_OCIGGS_SRC_USER_SEED_DATA.sql`.

 **Tip:**

You may need to run each statement separately for the SQL tool to execute the scripts successfully.

- e. To verify that the tables were created successfully, close the SQL window and reopen it again. In the Navigator tab, look for the `SRC_OCIGLL` schema and then select tables from their respective dropdowns.
- f. Enable supplemental logging:
 - i. Clear the SQL Worksheet.
 - ii. Enter the following statement, and then click **Run Statement**:

```
ALTER PLUGGABLE DATABASE ADD SUPPLEMENTAL LOG DATA;
```

3. Set up the target Autonomous Data Warehouse:
 - a. In the Oracle Cloud console, select your ADW instance from the Autonomous Databases page to view its details and access DB tools.
 - b. Click **Database Actions**.
 - c. In the Database Actions menu, under Development, select **SQL**.
 - d. Copy and paste the script from previously downloaded `OCIGLL_OCIGGS_SETUP_USERS_ADW.sql` into the SQL worksheet.
 - e. Click **Run Script**. The Script Output tab displays confirmation messages.
 - f. Clear the SQL worksheet and then copy and paste the SQL script from `OCIGLL_OCIGGS_SRC_MIRROR_USER_SEED_DATA.sql`
 - g. Click **Run Script**.

Task 1: Create OCI GoldenGate resources

This quickstart example requires deployments and connections for both the source and target.

1. Create an Oracle deployment for the source Autonomous Transaction Processing instance.
2. Create a Big Data deployment for the target Autonomous Data Warehouse.
3. Create a connection for the source Autonomous Transaction Processing instance.
4. Create a connection for the target Autonomous Data Warehouse instance.
5. Create connection for OCI Object Storage.
6. (Optional) If your Big Data deployment doesn't have a public endpoint, then create a connection to GoldenGate, and then assign this connection to the source Oracle deployment.
7. Assign the ATP connection to the source Oracle deployment.
8. Assign the ADW connection the target Big Data deployment.

9. Assign the OCI Object Storage connection to the target Big Data deployment.

Task 2: Add the Extract

1. On the Deployments page, select the source Autonomous Transaction Processing deployment.
2. On the deployment details page, click **Launch Console**.
3. Log in with the source deployment's administrator username and password.
4. Add transaction information.
5. Add an Extract.

Task 3: Add and run a Distribution Path

1. Create a user for the Distribution Path in the target Big Data deployment.
 - a. On the Deployments page, select the target deployment to view its details.
 - b. On the deployment details page, click **Launch Console**.
 - c. Sign in to the target OCI GoldenGate deployment console using the Administrator username and password.
 - d. In the OCI GoldenGate deployment console, open the navigation menu, and then click **Administrator**.
 - e. On the Users page, click **Add New User** (plus icon).
 - f. Complete the fields as follows, and then click **Submit**.
 - For **Username**, enter a name, such as `ggsnet`.
 - From the **Role** dropdown, select **Operator**.
 - Enter a password twice for verification.

The new user appears in the Users list.

2. In the source ATP deployment console, add a credential for the user created in Step 1.
 - a. In the source ATP deployment console, click **Administration Service**, open the navigation menu, and then select **Configuration**.
 - b. On the Credentials page, click **Add Credential**, and then complete the fields as follows:
 - For **Credential Domain**, enter `GGSNetwork`.
 - For **Credential Alias**, enter `dpuser`
 - For **User ID**, enter the name of the user created in Step 1 (`ggsnet`)
 - Enter the user's password twice for verification.
 - c. Click **Submit**.

The credential appears in the Credentials list.

3. Click **Distribution Service**, and then click **Add Path** (plus icon).
4. Complete the Add Path form fields as follows, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.

- b. For **Source Extract**, select the Extract created in Task 2.
- c. For **Source Trail Name**, select the Extract Trail from Task 2.
- d. For **Target Authentication Method**, select **UserID Alias**.
- e. For **Target** protocol, select **wss**.
- f. For **Target Host**, enter the host domain of the target deployment.

 **Note:**

You can copy and paste the URL of the target OCI GoldenGate deployment console and remove the https:// and any trailing slashes.

- g. For **Port Number**, enter 443.
- h. For **Trail Name**, enter a two-character name, such as E1.
- i. For **Domain**, enter the name of the Credential Domain created in Step 2 (GGSNetwork).
- j. For **Alias**, enter the Credential Alias created in Step 2 (dpuser).

You're returned to the Distribution Service Overview page where you can view the status of the created path.

5. In the target Big Data deployment console, review the Receiver Path.
 - a. In the target Big Data OCI GoldenGate deployment console, click **Receiver Service**.
 - b. Review the path details. This path was created as a result of the Distribution Path created in the previous step.

Task 4: Add and run the Replicat

1. In the target Big Data deployment console, click **Administrator Service**, and then click **Add Replicat** (plus icon).
2. On the Add Replicat page, under Replicat type, select **Classic Replicat**, and then click **Next**.
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name.
 - b. For **Trail Name**, enter the name of the Trail from Task 2.
 - c. For **Target**, select **Autonomous Data Warehouse**.
 - d. For **Alias**, select the **OCI Object Storage connection** and the **Autonomous Data Warehouse connection** created in Task 1.
 - e. For **Via Dependent Alias**, select the OCI Object Storage connection created in Task 1
4. On the Replicat Parameters page, change the MAP line to the following, and then click **Next**:

```
MAP SRC_OCIGLL.*, TARGET SRCMIRROR_OCIGLL.*;
```

5. On the Properties page, configure the following properties:
 - a. `gg.eventhandler.oci.compartmentID`: Add the OCID of the compartment in which the OCI Object Storage bucket is stored.
 - b. `gg.eventhandler.oci.bucketMappingTemplate`: Add the name of the OCI Object Storage bucket.
6. Click **Create and Run**.

You're returned to the Overview page, where you can review the Replicat details.

Task 5: Verify the replication

1. In the Oracle Cloud console, from the navigation menu, select Oracle Database, and then select **Autonomous Transaction Processing**.
2. In the list of Autonomous Transaction Processing instances, select your source instance to view its details.
3. On the database details page, click **Database Actions**.



Note:

You should be automatically logged in. If not, log in with the database credentials.

4. On the Database Actions home page, select **SQL**.
5. Enter the following into the worksheet and click **Run Script**.
6. In the source GoldenGate OCI GoldenGate deployment console, select the Extract name, and then click **Statistics**. Verify that `SRC_OCIGLL.SRC_CUSTOMER` has 7 inserts.

```

Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
  values (1001,0,'Brendt','Paul','10 Jasper Blvd.',107,'(212) 555
2146',19,10);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
  values (1002,0,'McCarthy','Robin','27 Pasadena Drive',11,'(214) 555
3075',29,11);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
  values (1003,0,'Travis','Peter','7835 Hartford Drive',12,'(510) 555
4448',34,12);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
  values (1004,0,'Larson','Joe','87 Carmel Blvd.',13,'(213) 555
5095',45,13);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
  values (1005,0,'Goldschmidt','Tony','91 Torre drive',14,'(619) 555
6529',55,20);
Insert into SRC_OCIGLL.SRC_CUSTOMER
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_PERS_ID)
  values (1006,0,'Baker','William','2890 Grant Avenue',15,'(312) 555

```

```
7040',64,21);  
Insert into SRC_OCIGLL.SRC_CUSTOMER  
(CUSTID,DEAR,LAST_NAME,FIRST_NAME,ADDRESS,CITY_ID,PHONE,AGE,SALES_P  
RS_ID) values (1007,0,'Swenson','Jack','64 Imagination  
Drive',19,'(202) 555 8125',74,22);
```

7. In the target Big Data OCI GoldenGate deployment console, select the Replicat name, and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CUSTOMER has 7 inserts.
8. In target Autonomous Data Warehouse Cloud SQL console, execute the following command to validate the data replicated:

```
select * from SRCMIRROR_OCIGLL.SRC_CUSTOMER;
```

Task 6: Monitor and maintain processes

1. Monitor the replication process.
2. Manage Trail files.

Replicate data from MongoDB to Autonomous JSON Database

Learn to replicate data from MongoDB to Autonomous JSON Database using OCI GoldenGate.

Oracle Autonomous JSON Database is a cloud document database service that makes it simple to develop JSON-centric applications. It features NoSQL-style document APIs (Oracle SODA and Oracle Database API for MongoDB), serverless scaling, high performance ACID transactions, comprehensive security, and low pay-per-use pricing. Learn more about [Autonomous JSON Database](#).

The OCI GoldenGate Big Data deployment type supports no down-time migrations from MongoDB to Autonomous JSON Database. OCI GoldenGate supports both Initial Load Extract and Change Data Capture (CDC) extract from MongoDB.

This quickstart details the process to configure OCI GoldenGate for no down-time migrations from MongoDB to Autonomous JSON Database.

Replicate data from PostgreSQL to Snowflake

Learn to replicate data from PostgreSQL to Snowflake using OCI GoldenGate.

Before you begin

To successfully complete this quickstart, you must have the following:

- [A PostgreSQL installation](#) to serve as the source database (Installation instructions follow in Task 0).
- Open port 5432 in your VCN's security list.
- [A Snowflake database](#) to serve as the target database.

Set up the environment

To set up the environment for this Quickstart:

1. Install PostgreSQL.

a. Install PostgreSQL server:

```
sudo yum install postgresql-server
```

b. Install postgresql-contrib module to avoid [this SQL exception](#):

```
sudo yum install postgresql-contrib
```

c. Create a new PostgreSQL database cluster:

```
sudo postgresql-setup --initdb
```

d. Enable the postgresql.service:

```
sudo systemctl enable postgresql.service
```

e. Start the postgresql.service:

```
sudo systemctl start postgresql.service
```

2. By default, PostgreSQL only allows local connections. [Allow remote connectivity to PostgreSQL](#).

a. In `/var/lib/pgsql/data/postgresql.conf`, prepare the database for replication:

i. Locate and uncomment `listen_addresses = 'localhost'` and change `localhost` to an asterisk (*):

```
listen_addresses = '*'
```

ii. Set the following parameters as follows:

- `wal_level = logical`
- `max_replication_slots = 1`
- `max_wal_senders = 1`
- `track_commit_timestamp = on`

Note:

Configure `/var/lib/pgsql/data/pg_hba.conf` to ensure that client authentication is set to allow connections from an Oracle GoldenGate host. For example, add the following:

```
#Allow connections from remote hosts
host    all    all    0.0.0.0/0    md5
```

See [The pg_hba.conf File](#) for more information.

- b. Restart PostgreSQL server:

```
sudo systemctl restart postgresql.service
```

3. If using Oracle Cloud Compute to host PostgreSQL, open port 5432:

```
sudo firewall-cmd --permanent --add-port=5432/tcp
sudo firewall-cmd --reload
sudo firewall-cmd --list-all
```

4. Open port 5432 in your VCN's security list.

5. [Connect to PostgreSQL](#).

```
> sudo su - postgres
> psql
```

6. Set up PostgreSQL.

- a. Download and run [seedSRCOCIGLL_PostgreSQL.sql](#) to set up the database and load the sample data.

- b. Run the following commands to set up the user:

```
create user ggadmin with password 'W3lcome@1234';
alter user ggadmin with SUPERUSER;
GRANT ALL PRIVILEGES ON DATABASE ociggl TO ggadmin;
```

7. Set up Snowflake:

- a. [Create a GoldenGate user in Snowflake](#) with appropriate privileges.
- b. Create target tables using [sample schema](#).
- c. Ensure the tables and user have been successfully created.

Task 1: Create the OCI GoldenGate Resources

This quickstart example requires deployments and connections for both the source and target.

1. Create a deployment for the source PostgreSQL database.
2. Create a Big Data deployment for the target Snowflake database.
3. Create a PostgreSQL connection with the following values:
 - a. For **Type**, select **PostgreSQL Server** from the dropdown.
 - b. For **Database name**, enter `ociggl`.
 - c. For **Host**, enter the public IP of the Compute instance that PostgreSQL runs on.
 - d. For **Port**, enter 5432.
 - e. For **Username**, enter `ggadmin`.
 - f. For **Password**, enter `w3lcome@1234`.
 - g. For **Security Protocol**, select **Plain** from the dropdown.
4. Create a Snowflake connection with the following values:

- a. For **Connection URL**, enter `jdbc:snowflake://<account_identifier>.snowflakecomputing.com/?warehouse=<warehouse name>&db=OCIGLL`.

 **Note:**

Ensure you replace `<account_identifier>` and `<warehouse name>` with the appropriate values.

- b. For **Authentication Type**, select **Basic authentication** from the dropdown.
 - c. For **Username**, enter a name.
 - d. For **Password**, enter a password.
5. (Optional) If your Big Data deployment doesn't have a public endpoint, then create a connection to GoldenGate, and then assign this connection to the source PostgreSQL deployment.
 6. Assign the source PostgreSQL connection to the PostgreSQL deployment.
 7. Assign the Snowflake connection to the target Big Data deployment.

Task 2: Create the Extracts

1. Enable supplemental logging:
 - a. Launch the PostgreSQL GoldenGate deployment console:
 - i. From the Deployments page, select the PostgreSQL deployment to view its details.
 - ii. On the PostgreSQL deployment details page, click **Launch console**.
 - iii. On the deployment console sign in page, enter the GoldenGate admin credentials provided in Task 1, step 1.
 - b. After signing in, open the navigation menu, and then click **Configuration**.
 - c. Click **Connect**. Checkpoint table and TRANDATA fields appear if the connection is successful.
 - d. Next to TRANDATA Information, click **Add TRANDATA** (plus icon).
 - e. For **Table Name**, enter `src_ociggl1.*;`, and then click **Submit**.

 **Note:**

You only need to click Submit once. Use the search field to search for `src_ociggl1` and verify the tables were added.

2. Add the Change Data Capture Extract:
 - a. From the navigation menu, click **Overview**.
 - b. On the Administration Service page, click **Add Extract** (plus icon).
 - c. On the Extract Type page, select **Change Data Capture Extract**, and then click **Next**.
 - d. Complete the Extract Options as follows, and then click **Next**:

- i. For **Process Name**, enter a name for the Extract, such as `ECDPCSQL`.
 - ii. For **Credential Domain**, select **Oracle GoldenGate**.
 - iii. For **Credential Alias**, select the alias.
 - iv. For **Begin**, select **Now**.
 - v. For **Trail Name**, enter a two-character trail name, such as `P1`.
- e. On the Extract Parameters page, add the following:

```
TABLE SRC_OCIGGLL.*;
```

- f. Click **Create and Run**.
3. Add the Initial Load Extract:
- a. On the Administration Service Overview page, click **Add Extract** (plus icon).
 - b. On the Extract Type page, select **Initial Load Extract**, and then click **Next**.
 - c. Complete the Extract Options as follows, and then click **Next**:
 - i. For **Process Name**, enter a name for the Extract, such as `EINIPSQL`.
 - ii. For **Credential Domain**, select **Oracle GoldenGate**.
 - iii. For **Credential Alias**, select the alias.
 - iv. For **Trail Name**, enter a two-character trail name, such as `I1`.
 - d. In the Extract Parameters text area, add the following:

```
EXTRACT EINIPSQL
USERIDALIAS PostgreSQL_Compute, DOMAIN OracleGoldenGate
EXTFILE I1, PURGE
TABLE src_ociggll.*;
```

 **Note:**

Ensure that you remove the `SOURCEDB` parameter in front of `USERIDALIAS` before you move on.

- e. Click **Create and Run**.
You return to the Administration Service Overview page, where you can observe the Extract starting.

Task 3: Create the Distribution Path for Initial Load Extract

To create a Distribution Path for Initial Load Extract, complete the following:

1. In the Oracle Cloud console, on the Deployments page, select the target Big Data deployment.
2. On the deployment details page, click **Launch Console**. Log in with the admin user details created in task 1, step 2.
3. Create a user for the Distribution Path.
 - a. Open the navigation menu, and then click **Administrator**.

- b. Click **Add New User** (plus icon), complete the fields as follows, and then click **Submit**:
 - For **Username**, enter `ggsnet`.
 - For **Role**, select **Operator**.
 - Enter the password twice for verification.
4. In the source PostgreSQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click **Add Credential** (plus icon), complete the fields as follows, and then click **Submit**:
 - For **Credential Domain**, enter `GGNetwork`.
 - For **Credential Alias**, enter `dpuser`.
 - For **Database Name**, you can enter any name and leave the Database Server and Port fields blank, or use the default values.
 - For **User ID**, enter `ggsnet`
 - For **Password**, enter the same password used in the previous step.
5. In the source PostgreSQL deployment console, create a Distribution Path.
 - a. Click **Distribution Service**, and then click **Add Path** (plus icon).
 - b. Complete the following fields, and click **Create and Run**:
 - i. For **Path Name**, enter a name for this path.
 - ii. For **Source Trail**, leave blank.
 - iii. For **Trail Name**, enter the Initial Load Extract trail name (`I1`).
 - iv. For **Target Authentication Method**, select **UserID Alias**.
 - v. For **Target**, select `wss`.
 - vi. For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the `https://` or any trailing slashes.
 - vii. For **Port Number**, enter `443`.
 - viii. For **Trail Name**, enter `I1`.
 - ix. For **Domain**, enter the domain name created in the previous step.
 - x. For **Alias**, enter the alias created in the previous step.

You return to the Distribution Service Overview page where you can review the path created.
6. In the target Big Data deployment console, review the Receiver path created as a result of the Distribution path:
 - a. Click **Receiver Service**.
 - b. Review the Receiver path details.

Task 4: Add the Replicat for Initial Load

1. In the target Big Data deployment console, add a checkpoint table.

- a. Open the navigation menu, and then select **Configuration**.
 - b. Click the connect icon for the target Snowflake database.
 - c. Click **Add Checkpoint** (plus icon).
 - d. For **Checkpoint Table**, enter `SRCMIRROR_OCIGLL.CHECKTABLE`.
 - e. Click **Submit**.
2. Add the Initial Load Replicat.
 - a. In the navigation menu, click **Overview**, and then click **Add Replicat** (plus icon).
 - b. On the Add Replicat page, under Replicat type, select either **Classic** or **Coordinated**, and then click **Next**.
 - c. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - i. For **Process Name**, enter a name, such as `RIL`.
 - ii. For **Credential Domain**, select the domain for the Snowflake connection.
 - iii. For **Credential Alias**, select the alias of the Snowflake connection.
 - iv. For **Trail Name**, enter the name of the Trail from Task 2 (`I1`).
 - v. For **Target**, select the target **Snowflake** connection from the dropdown.
 - vi. For **Available aliases**, select an alias from the dropdown, such as `Snowflake`.
 - vii. (Optional) **Enable external storage** to select an available staging location from the dropdown.
 - d. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
INSERTALLRECORDS
MAP src_ocigll.src_city, TARGET SRCMIRROR_OCIGLL.SRC_CITY;
MAP src_ocigll.src_region, TARGET SRCMIRROR_OCIGLL.SRC_REGION;
MAP src_ocigll.src_customer, TARGET
SRCMIRROR_OCIGLL.SRC_CUSTOMER;
MAP src_ocigll.src_orders, TARGET SRCMIRROR_OCIGLL.SRC_ORDERS;
MAP src_ocigll.src_order_lines, TARGET
SRCMIRROR_OCIGLL.SRC_ORDER_LINES;
MAP src_ocigll.src_product, TARGET
SRCMIRROR_OCIGLL.SRC_PRODUCT;
```

- e. On the Properties page, review the properties, and then click **Create and Run**.
You return to the Overview page, where you can review the Replicat details.
3. To verify the Initial Load, connect to Snowflake database and run following queries:

```
select * from SRCMIRROR_OCIGLL.SRC_CITY;
select * from SRCMIRROR_OCIGLL.SRC_CUSTOMER;
```

The output should return the data that was loaded into the target database tables as a result of the Initial Load.

Task 5: Create the Distribution Path for Change Data Capture

To create a Distribution Path for Change Data Capture, complete the following:

1. In the source PostgreSQL deployment console, click **Distribution Service**.
2. On the Paths page, click **Add Path**.
3. On the Add Path page, complete the following fields, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.
 - b. For **Source Extract**, select the Change Data Capture Extract (E CDCPSQL).
 - c. For **Trail Name**, select the Change Data Capture Extract trail file (P1).
 - d. For **Target Authentication Method**, select **UserID Alias**.
 - e. For **Target**, select **wss**.
 - f. For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the https:// or any trailing slashes.
 - g. For **Port Number**, enter 443.
 - h. For **Trail Name**, enter P1.
 - i. For **Domain**, enter the domain name created in task 3.
 - j. For **Alias**, enter the alias created in task 3.
4. In the target Big Data deployment console, click **Receiver Service**, and then review the Receiver path created.

Task 6: Add a Replicat for Change Data Capture

Perform updates to the source PostgreSQL database to verify replication to Snowflake.

1. Add the Change Data Replicat.
 - a. In the target Big Data deployment console, click **Administrator Service**, and then click **Add Replicat** (plus icon).
 - b. On the Add Replicat page, under Replicat type, select either **Classic** or **Coordinated**, and then click **Next**.
 - c. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - i. For **Process Name**, enter a name, such as RCDC.
 - ii. For **Credential Domain**, select the domain for the Snowflake connection.
 - iii. For **Credential Alias**, select the alias of the Snowflake connection.
 - iv. For **Trail Name**, enter the name of the Trail from Task 2 (P1).
 - v. For **Target**, select the target **Snowflake** connection from the dropdown.
 - vi. For **Available aliases**, select an alias from the dropdown, such as Snowflake.
 - vii. (Optional) **Enable external storage** to select an available staging location from the dropdown.

- d. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP src_ociggl.src_city, TARGET SRCMIRROR_OCIGLL.SRC_CITY;
MAP src_ociggl.src_region, TARGET SRCMIRROR_OCIGLL.SRC_REGION;
MAP src_ociggl.src_customer, TARGET
SRCMIRROR_OCIGLL.SRC_CUSTOMER;
MAP src_ociggl.src_orders, TARGET SRCMIRROR_OCIGLL.SRC_ORDERS;
MAP src_ociggl.src_order_lines, TARGET
SRCMIRROR_OCIGLL.SRC_ORDER_LINES;
MAP src_ociggl.src_product, TARGET
SRCMIRROR_OCIGLL.SRC_PRODUCT;
```

- e. On the Properties page, review the properties, and then click **Create and Run**.

You return to the Overview page, where you can review the Replicat details.

2. Verify the Change Data Capture:

- a. Perform updates to the source PostgreSQL database to verify replication to Snowflake. Run the following script to perform inserts into the PostgreSQL database:

```
Insert into src_ociggl.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1000,'Houston',20,743113);
Insert into src_ociggl.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1001,'Dallas',20,822416);
Insert into src_ociggl.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1002,'San
Francisco',21,157574);
Insert into src_ociggl.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1003,'Los
Angeles',21,743878);
Insert into src_ociggl.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1004,'San
Diego',21,840689);
Insert into src_ociggl.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1005,'Chicago',23,616472);
Insert into src_ociggl.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1006,'Memphis',23,580075);
Insert into src_ociggl.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1007,'New York
City',22,124434);
Insert into src_ociggl.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1008,'Boston',22,275581);
Insert into src_ociggl.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1009,'Washington
D.C.',22,688002);
```


- b. In the source PostgreSQL deployment console, select the Change Data Capture Extract name (RCDC), and then click **Statistics**. Verify that `src_ocigg11.src_city` has 10 inserts.

 **Note:**

If the Extract captured no inserts, then restart the ECDCPSQL Extract.

- c. In the target Big Data deployment console, select the Change Data Capture Replicat name (RCDC), review its **Details** and **Statistics** to verify the number of Inserts.

Task 7: Monitor and maintain processes

1. Monitor performance.
2. Manage Trail files.

Send Data from OCI MySQL Heatwave to Azure Event Hubs

Learn how to use OCI GoldenGate to replicate data from OCI MySQL Heatwave to Azure Event Hubs.

Before you begin

To successfully complete this quickstart, you must have the following:

- OCI Bastion, to connect to OCI MySQL Heatwave Database, load the sample database, and perform inserts.
- A OCI MySQL Heatwave Database to serve as the source database.
- [An Azure Event Hubs namespace and an event hub created](#)

 **Note:**

Kafka Surface is not enabled in Azure Event Hubs Basic Tier. Standard or Premium tier is required.

Task 0: Set up the environment

To set up the environment for this Quickstart:

1. [Download the sample data script](#), and then run the script on the OCI MySQL Heatwave database to create the database and load the data.
2. Create a `ggadmin` user using the following script. Remember to replace `<ggadmin-password>` with a valid password:

```
CREATE USER 'ggadmin' IDENTIFIED BY '<ggadmin-password>';  
GRANT SELECT, REPLICATION SLAVE, REPLICATION CLIENT, CREATE,CREATE VIEW,  
EVENT, INSERT, UPDATE, DROP,EXECUTE, DELETE ON *.* TO 'ggadmin';
```

3. Open port 3306, through which OCI GoldenGate can connect.

- a. In the Oracle Cloud console, locate the subnet that the OCI MySQL Heatwave database uses.
- b. In the security list of the subnet, create an Ingress rule for port TCP/3306.

Task 1: Create the OCI GoldenGate resources

1. Create a MySQL deployment for the source OCI MySQL Heatwave database.
2. Create a Big Data deployment for the target Azure Event Hubs.
3. Create a connection to the source MySQL Heatwave Database.
4. Create an Azure Event Hubs connection.
5. Create a connection to GoldenGate, and then assign this connection to the source MySQL deployment.
6. Assign the source connection to the source MySQL deployment.
7. Assign the target connection to the target Big Data deployment.

Task 2: Create the Extracts

1. On the Deployments page, select the MySQL deployment created in Task 1.
2. On the deployment details page, click **Launch Console**.
3. Sign in to the source MySQL deployment console using the Administrator credentials specified when you created the deployment in Task 1.
4. Open the navigation menu, and then select **Configuration**.
5. Review the credentials for the MySQL Database Server connection. Take note of the **Domain** and **Alias**.
6. Return to the Administration Service Overview page.
7. On the Administration Service Overview page, click **Add Extract (plus icon)**.
8. On the Extract Type page, select **Change Data Capture Extract**, and then click **Next**.
9. Complete the Extract Options as follows, and then click **Next**:
 - a. For **Process Name**, enter a name for the Extract, such as CDCEXT.
 - b. For **Credential Domain**, select the domain from Step 5.
 - c. For **Credential Alias**, select the alias from Step 5.
 - d. For **Trail Name**, enter a two-character trail name, such as C1.
10. Enable **Remote** if capturing from a MySQL Database that is not using global transaction identifiers (GTIDs).
11. On the Extract Parameters page, add the following:

```
TABLE SRC_OCIGLL.*;
```

12. Click **Create and Run**.

You're returned to the Administration Service Overview page, where you can observe the Extract starting.

Task 3: Create the Distribution Path to target Big Data deployment

1. On the Deployments page, select the target Big Data deployment.
2. On the deployment details page, click **Launch Console**.
3. Create a user for the Distribution Path:
 - a. Open the navigation menu, and then click **Administrator**.
 - b. Click **Add New User (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Username**, enter `ggsnet`.
 - For **Role**, select **Operator**.
 - Enter the password twice for verification.
4. In the source MySQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click **Add Credential (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Credential Domain**, enter `GGSNetwork`.
 - For **Credential Alias**, enter `dpuser`.
 - For **Database Name**, you can enter any name and leave the Database Server and Port fields blank, or use the default values.
 - For **User ID**, enter `ggsnet`.
 - For **Password**, enter the same password used in the previous step.
5. Create a Distribution Path:
 - a. In the source MySQL deployment console, click **Distribution Service**, and then click **Add Path (plus icon)**.
 - b. Complete the following fields, and click **Create and Run**:
 - For **Path Name**, enter a name for this path.
 - For **Source Extract**, select the **Change Data Capture Extract (CDCEXT)**.
 - For **Trail Name**, select the Change Data Capture Extract trail file (C1).
 - For **Target Authentication Method**, select **UserID Alias**.
 - For **Target**, select **wss**.
 - For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the `https://` or any trailing slashes.
 - For **Port Number**, enter **443**.
 - For **Trail Name**, enter `C1`.
 - For **Domain**, enter the domain name created in the previous step.
 - For **Alias**, enter the alias created in the previous step.

You're returned to the Distribution Service Overview page where you can review the path created.

6. In the target OCI GoldenGate Big Data deployment console, review the Receiver path created as a result of the Distribution path:
 - a. Click **Receiver Service**.
 - b. Review the Receiver path details.

Task 4: Add a Replicat for Azure Event Hubs

1. Add a Replicat:
 - a. In Big Data deployment click **Administrator Service**, and then click **Add Replicat (plus icon)**.
 - b. Select **Classic Replicat** and then click **Next**.
2. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name.
 - b. For **Trail Name**, enter the name of the Trail from Task 2.
 - c. For **Target**, select **Azure Event Hubs**.
 - d. For **Alias**, select the Azure Event Hubs connection created in Task 1.
3. On the Replicat Parameters page, leave the default, and then click **Next**.
4. On the Properties page, provide a topic name for `topicMappingTemplate`.

 **Note:**

The topic name can be a static name or a template keyword for a dynamic topic name. If a topic is not present, OCI GoldenGate generates one for you.

5. Click **Create and Run**.

Task 5: Verify Data Replication from OCI MySQL Heatwave to Azure Event Hubs

Perform updates to the source OCI MySQL Heatwave database to verify replication to Azure Event Hubs.

1. In OCI Bastion, create an SSH port forwarding session using MySQL IP and `port 3306`. Add your public SSH key.
2. Connect to MySQL in Cloud Shell using your private key and `port 3306`.
3. After connecting successfully, run the following command:

```
mysqlsh admin@localhost:3306 --sql
```

4. Run the following script to perform inserts into the OCI MySQL Heatwave database:

```
use SRC_OCIGLL;
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1000,'Houston',20,743113);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
```

```
(1001, 'Dallas', 20, 822416);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1002, 'San Francisco', 21, 157574);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1003, 'Los Angeles', 21, 743878);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1004, 'San Diego', 21, 840689);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1005, 'Chicago', 23, 616472);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1006, 'Memphis', 23, 580075);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1007, 'New York City', 22, 124434);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1008, 'Boston', 22, 275581);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1009, 'Washington D.C.', 22, 688002);
commit;
```

5. In the source MySQL deployment console, select the Change Data Capture Extract name (CDCEXT), and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CITY has 10 inserts.
6. In target Big Data deployment console, select the Replicat Name and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CITY has 10 inserts.
7. In Azure Event Hubs, you can check the message statistics.

Replicate Data from OCI MySQL Heatwave to Amazon Kinesis

Learn how to use OCI GoldenGate to replicate data from OCI MySQL Heatwave to Amazon Kinesis.

Before you begin

To successfully complete this quickstart, you must have the following:

- OCI Bastion, to connect to OCI MySQL Heatwave Database, load the sample database, and perform inserts.
- A OCI MySQL Heatwave Database to serve as the source database.
- [Amazon S3 Access Key & Secret](#)
- [IAM Policies for Amazon Kinesis](#)

Task 0: Set up the environment

To set up the environment for this Quickstart:

1. [Download the sample data script](#), and then run the script on the OCI MySQL Heatwave database to create the database and load the data.
2. Create a `ggadmin` user using the following script. Remember to replace `<ggadmin-password>` with a valid password:

```
CREATE USER 'ggadmin' IDENTIFIED BY '<ggadmin-password>';
GRANT SELECT, REPLICATION SLAVE, REPLICATION CLIENT, CREATE,CREATE VIEW,
EVENT, INSERT, UPDATE, DROP,EXECUTE, DELETE ON *.* TO 'ggadmin';
```

3. Open `port 3306`, through which OCI GoldenGate can connect.
 - a. In the Oracle Cloud console, locate the subnet that the OCI MySQL Heatwave database uses.
 - b. In the security list of the subnet, create an Ingress rule for `port TCP/3306`.

Task 1: Create the OCI GoldenGate resources

1. Create a MySQL deployment for the source OCI MySQL Heatwave database.
2. Create a Big Data deployment for Amazon Kinesis target.
3. Create a connection to the source MySQL Heatwave Database.
4. Connect to Amazon Kinesis.
5. Create a connection to GoldenGate, and then assign this connection to the source MySQL deployment.
6. Assign the source connection to the source MySQL deployment.
7. Assign the target connection to the target Big Data deployment.

Task 2: Create the Extracts

1. On the Deployments page, select the MySQL deployment created in Task 1.
2. On the deployment details page, click **Launch Console**.
3. Sign in to the source MySQL deployment console using the Administrator credentials specified when you created the deployment in Task 1.
4. Open the navigation menu, and then select **Configuration**.
5. Review the credentials for the MySQL Heatwave connection. Take note of the **Domain** and **Alias**.
6. Return to the Administration Service Overview page.
7. On the Administration Service Overview page, click **Add Extract (plus icon)**.
8. On the Extract Type page, select **Change Data Capture Extract**, and then click **Next**.
9. Complete the Extract Options as follows, and then click **Next**:
 - a. For **Process Name**, enter a name for the Extract, such as `CDCEXT`.
 - b. For **Credential Domain**, select the domain from Step 5.
 - c. For **Credential Alias**, select the alias from Step 5.
 - d. For **Trail Name**, enter a two-character trail name, such as `c1`.
10. Enable **Remote** if capturing from a MySQL Database that is not using global transaction identifiers (GTIDs).
11. On the Extract Parameters page, add the following:

```
TABLE SRC_OCIGLL.*;
```

12. Click **Create and Run**.

You're returned to the Administration Service Overview page, where you can observe the Extract starting.

Task 3: Create the Distribution Path to target Big Data deployment

1. On the Deployments page, select the target Big Data deployment.
2. On the deployment details page, click **Launch Console**.
3. Create a user for the Distribution Path:
 - a. Open the navigation menu, and then click **Administrator**.
 - b. Click **Add New User (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Username**, enter `ggsnet`.
 - For **Role**, select **Operator**.
 - Enter the password twice for verification.
4. In the source MySQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click **Add Credential (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Credential Domain**, enter `GGSNetwork`.
 - For **Credential Alias**, enter `dpuser`.
 - For **Database Name**, you can enter any name and leave the Database Server and Port fields blank, or use the default values.
 - For **User ID**, enter `ggsnet`.
 - For **Password**, enter the same password used in the previous step.
5. Create a Distribution Path:
 - a. In the source MySQL deployment console, click **Distribution Service**, and then click **Add Path (plus icon)**.
 - b. Complete the following fields, and click **Create and Run**:
 - For **Path Name**, enter a name for this path.
 - For **Source Extract**, select the **Change Data Capture Extract (CDCEXT)**.
 - For **Trail Name**, select the Change Data Capture Extract trail file (C1).
 - For **Target Authentication Method**, select **UserID Alias**.
 - For **Target**, select **wss**.
 - For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the `https://` or any trailing slashes.
 - For **Port Number**, enter **443**.
 - For **Trail Name**, enter `C1`.
 - For **Domain**, enter the domain name created in the previous step.
 - For **Alias**, enter the alias created in the previous step.

You're returned to the Distribution Service Overview page where you can review the path created.

6. In the target OCI GoldenGate Big Data deployment console, review the Receiver path created as a result of the Distribution path:
 - a. Click **Receiver Service**.
 - b. Review the Receiver path details.

Task 4: Add a Replicat for Amazon Kinesis

1. Add a Replicat:
 - a. In Big Data deployment click **Administrator Service**, and then click **Add Replicat (plus icon)**.
 - b. Select **Classic Replicat** and then click **Next**.
2. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name.
 - b. For **Trail Name**, enter the name of the Trail from Task 2.
 - c. For **Target**, select **Amazon Kinesis**.
 - d. For **Alias**, select the Amazon Kinesis connection created in Task 1.
3. On the Replicat Parameters page, leave the default, and then click **Next**.
4. On the Properties page, edit the fields marked #TODO.
 - a. `gg.handler.kinesis.region`: provide the Amazon Web Services (AWS) region for the target Kinesis stream.
 - b. `gg.handler.kinesis.streamMappingTemplate`: by default, it is set to `{tableName}` which will map the streams based on source table name. If you want to map to an existing data stream, you can provide static stream names or you can use Template Keywords to assign stream names dynamically.
5. Click **Create and Run**.

Task 5: Verify Data Replication from OCI MySQL Heatwave to Amazon Kinesis

Perform updates to the source OCI MySQL Heatwave database to verify replication to Amazon Kinesis.

1. In OCI Bastion, create an SSH port forwarding session using MySQL IP and `port 3306`. Add your public SSH key.
2. Connect to MySQL in Cloud Shell using your private key and `port 3306`.
3. After connecting successfully, run the following command:

```
mysqlsh admin@localhost:3306 --sql
```

4. Run the following script to perform inserts into the OCI MySQL Heatwave database:

```
use SRC_OCIGLL;
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1000,'Houston',20,743113);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1001,'Dallas',20,822416);
```



```
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1002,'San Francisco',21,157574);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1003,'Los Angeles',21,743878);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1004,'San Diego',21,840689);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1005,'Chicago',23,616472);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1006,'Memphis',23,580075);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1007,'New York City',22,124434);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1008,'Boston',22,275581);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1009,'Washington D.C.',22,688002);
commit;
```

5. In the source MySQL deployment console, select the Change Data Capture Extract name (CDCEXT), and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CITY has 10 inserts.
6. In target Big Data deployment console, select the Replicat Name and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CITY has 10 inserts.
7. In Amazon Kinesis, you can check the messages in target data stream/data viewer.

Replicate Data from OCI MySQL Heatwave to Google Cloud Storage

Learn how to use OCI GoldenGate to replicate data from OCI MySQL Heatwave to Google Cloud Storage.

Before you begin

To successfully complete this quickstart, you must have the following:

- OCI Bastion, to connect to OCI MySQL Heatwave Database, load the sample database, and perform inserts.
- A OCI MySQL Heatwave Database to serve as the source database.
- Google Cloud Storage account and [Google Cloud Service Account Key](#)
- Google Cloud Storage Bucket and Object permissions

Task 0: Set up the environment

To set up the environment for this Quickstart:

1. [Download the sample data script](#), and then run the script on the OCI MySQL Heatwave database to create the database and load the data.
2. Create a `ggadmin` user using the following script. Remember to replace `<ggadmin-password>` with a valid password:

```
CREATE USER 'ggadmin' IDENTIFIED BY '<ggadmin-password>';
GRANT SELECT, REPLICATION SLAVE, REPLICATION CLIENT, CREATE,CREATE VIEW,
EVENT, INSERT, UPDATE, DROP,EXECUTE, DELETE ON *.* TO 'ggadmin';
```

3. Open `port 3306`, through which OCI GoldenGate can connect.
 - a. In the Oracle Cloud console, locate the subnet that the OCI MySQL Heatwave database uses.
 - b. In the security list of the subnet, create an Ingress rule for `port TCP/3306`.

Task 1: Create the OCI GoldenGate resources

1. Create a MySQL deployment for the source OCI MySQL Heatwave database.
2. Create a Big Data deployment for the target Google Cloud Storage.
3. Create a connection to the source MySQL Heatwave Database.
4. Connect to Google Cloud Storage.
5. Create a connection to GoldenGate, and then assign this connection to the source MySQL deployment.
6. Assign the source connection to the source MySQL deployment.
7. Assign the target connection to the target Big Data deployment.

Task 2: Create the Extracts

1. On the Deployments page, select the MySQL deployment created in Task 1.
2. On the deployment details page, click **Launch Console**.
3. Sign in to the source MySQL deployment console using the Administrator credentials specified when you created the deployment in Task 1.
4. Open the navigation menu, and then select **Configuration**.
5. Review the credentials for the MySQL Database Server connection. Take note of the **Domain** and **Alias**.
6. Return to the Administration Service Overview page.
7. On the Administration Service Overview page, click **Add Extract (plus icon)**.
8. On the Extract Type page, select **Change Data Capture Extract**, and then click **Next**.
9. Complete the Extract Options as follows, and then click **Next**:
 - a. For **Process Name**, enter a name for the Extract, such as `CDCEXT`.
 - b. For **Credential Domain**, select the domain from Step 5.
 - c. For **Credential Alias**, select the alias from Step 5.
 - d. For **Trail Name**, enter a two-character trail name, such as `C1`.
10. Enable **Remote** if capturing from a MySQL Database that is not using global transaction identifiers (GTIDs).
11. On the Extract Parameters page, add the following:

```
TABLE SRC_OCIGLL.*;
```

12. Click **Create and Run**.

You're returned to the Administration Service Overview page, where you can observe the Extract starting.

Task 3: Create the Distribution Path to target Big Data deployment

1. On the Deployments page, select the target Big Data deployment.
2. On the deployment details page, click **Launch Console**.
3. Create a user for the Distribution Path:
 - a. Open the navigation menu, and then click **Administrator**.
 - b. Click **Add New User (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Username**, enter `ggsnet`.
 - For **Role**, select **Operator**.
 - Enter the password twice for verification.
4. In the source MySQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click **Add Credential (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Credential Domain**, enter `GGSNetwork`.
 - For **Credential Alias**, enter `dpuser`.
 - For **Database Name**, you can enter any name and leave the Database Server and Port fields blank, or use the default values.
 - For **User ID**, enter `ggsnet`.
 - For **Password**, enter the same password used in the previous step.
5. Create a Distribution Path:
 - a. In the source MySQL deployment console, click **Distribution Service**, and then click **Add Path (plus icon)**.
 - b. Complete the following fields, and click **Create and Run**:
 - For **Path Name**, enter a name for this path.
 - For **Source Extract**, select the **Change Data Capture Extract (CDCEXT)**.
 - For **Trail Name**, select the Change Data Capture Extract trail file (C1).
 - For **Target Authentication Method**, select **UserID Alias**.
 - For **Target**, select **wss**.
 - For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the `https://` or any trailing slashes.
 - For **Port Number**, enter **443**.
 - For **Trail Name**, enter `C1`.
 - For **Domain**, enter the domain name created in the previous step.
 - For **Alias**, enter the alias created in the previous step.

You're returned to the Distribution Service Overview page where you can review the path created.

6. In the target OCI GoldenGate Big Data deployment console, review the Receiver path created as a result of the Distribution path:
 - a. Click **Receiver Service**.
 - b. Review the Receiver path details.

Task 4: Add a Replicat for Google Cloud Storage

1. Add a Replicat:
 - a. In Big Data deployment click **Administrator Service**, and then click **Add Replicat (plus icon)**.
 - b. Select **Classic Replicat** and then click **Next**.
2. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name.
 - b. For **Trail Name**, enter the name of the Trail from Task 2.
 - c. For **Target**, select **Google Cloud Storage**.
 - d. For **Alias**, select the Azure Event Hubs connection created in Task 1.
3. On the Replicat Parameters page, leave the default, and then click **Next**.
4. On the Properties page, look for `gg.eventhandler.gcs.bucketMappingTemplate=<gcs bucket>`, and replace `<gcs bucket>` with the name of your bucket.



Note:

If target bucket does not exist, OCI GoldenGate creates one for you.

5. Click **Create and Run**.

Task 5: Verify Data Replication from OCI MySQL Heatwave to Google Cloud Storage

Perform updates to the source OCI MySQL Heatwave database to verify replication to Google Cloud Storage.

1. In OCI Bastion, create an SSH port forwarding session using MySQL IP and `port 3306`. Add your public SSH key.
2. Connect to MySQL in Cloud Shell using your private key and `port 3306`.
3. After connecting successfully, run the following command:

```
mysqlsh admin@localhost:3306 --sql
```

4. Run the following script to perform inserts into the OCI MySQL Heatwave database:

```
use SRC_OCIGLL;
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1000,'Houston',20,743113);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
```

```
(1001, 'Dallas', 20, 822416);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1002, 'San Francisco', 21, 157574);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1003, 'Los Angeles', 21, 743878);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1004, 'San Diego', 21, 840689);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1005, 'Chicago', 23, 616472);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1006, 'Memphis', 23, 580075);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1007, 'New York City', 22, 124434);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1008, 'Boston', 22, 275581);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1009, 'Washington D.C.', 22, 688002);
commit;
```

5. In the source MySQL deployment console, select the Change Data Capture Extract name (CDCEXT), and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CITY has 10 inserts.
6. In target Big Data deployment console, select the Replicat Name and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CITY has 10 inserts.
7. In Google Cloud Storage, you can check the messages in target data stream/data viewer.

Replicate data from PostgreSQL to Google BigQuery

Learn how to use OCI GoldenGate to replicate data from PostgreSQL to Google BigQuery.

Before you begin

To successfully complete this quickstart, you must have the following:

- [A PostgreSQL installation](#) to serve as the source database (Installation instructions follow in Task 0.)
- Open port 5432 in your VCN's security list.
- Create a connection to Google Cloud Storage.

 **Note:**

Please ensure that GCS bucket and the BigQuery dataset exist in the same location/region.

- [Google Cloud Service Account Key](#).
- Google Cloud Platform BigQuery Permissions.

Task 0: Set up the environment

To set up the environment for this Quickstart:

1. [Install PostgreSQL](#).

a. Install PostgreSQL server:

```
sudo yum install postgresql-server
```

b. Install postgresql-contrib module to avoid [this SQL exception](#):

```
sudo yum install postgresql-contrib
```

c. Create a new PostgreSQL database cluster:

```
sudo postgresql-setup --initdb
```

d. Enable the postgresql.service:

```
sudo systemctl enable postgresql.service
```

e. Start the postgresql.service:

```
sudo systemctl start postgresql.service
```

2. By default, PostgreSQL only allows local connections. [Allow remote connectivity to PostgreSQL](#).

a. In `/var/lib/pgsql/data/postgresql.conf`, prepare the database for replication:

i. Locate and uncomment `listen_addresses = 'localhost'` and change `localhost` to an asterisk (*):

```
listen_addresses = '*'
```

ii. Set the following parameters as follows:

- `wal_level = logical`
- `max_replication_slots = 1`
- `max_wal_senders = 1`
- `track_commit_timestamp = on`

 **Note:**

Configure `/var/lib/pgsql/data/pg_hba.conf` to ensure that client authentication is set to allow connections from an Oracle GoldenGate host. For example, add the following:

```
#Allow connections from remote hosts
host    all    all    0.0.0.0/0    md5
```

See [The pg_hba.conf File](#) for more information.

- b. Restart PostgreSQL server:

```
sudo systemctl restart postgresql.service
```

3. If using Oracle Cloud Compute to host PostgreSQL, open port 5432:

```
sudo firewall-cmd --permanent --add-port=5432/tcp
sudo firewall-cmd --reload
sudo firewall-cmd --list-all
```

4. Open port 5432 in your VCN's security list.

5. [Connect to PostgreSQL](#).

```
> sudo su - postgres
> psql
```

6. Set up PostgreSQL.

- a. Download and run [seedSRCOCIGLL_PostgreSQL.sql](#) to set up the database and load the sample data.

- b. Run the following commands to set up the user:

```
create user ggadmin with password 'W3lcome@1234';
alter user ggadmin with SUPERUSER;
GRANT ALL PRIVILEGES ON DATABASE ocigll TO ggadmin;
```

Task 1: Create the OCI GoldenGate resources

1. Create a deployment for the source PostgreSQL database.
2. Create a Big Data deployment for the target Google BigQuery.
3. Create a connection to the target Google BigQuery.
4. Create a connection to the source PostgreSQL database.
 - a. For **Type**, ensure that you select PostgreSQL Server.
 - b. For **Database name**, enter `ocigll`.
 - c. For **Host**, enter the public IP of the Compute instance that PostgreSQL runs on.
 - d. For **Port**, enter 5432.
 - e. For **Username**, enter `ggadmin`.
 - f. For **Password**, enter a password.
 - g. For **Security Protocol**, select **Plain**.
5. Create a connection to GoldenGate, and then assign this connection to the source PostgreSQL deployment.
6. Assign the source connection to the source PostgreSQL deployment..
7. Assign the target connection to the target Big Data deployment.
8. Enable supplemental logging:
 - a. Launch the PostgreSQL GoldenGate deployment console:

- i. From the Deployments page, select the PostgreSQL deployment to view its details.
- ii. On the PostgreSQL deployment details page, click **Launch console**.
- iii. On the deployment console sign in page, enter the GoldenGate admin credentials provided in Step 1.
- b. After signing in, open the navigation menu, and then click **Configuration**.
- c. Click **Connect**. Checkpoint table and TRANDATA fields appear if the connection is successful.
- d. Next to TRANDATA Information, click **Add TRANDATA** (plus icon).
- e. For Table Name, enter `src_ociggl1.*;`, and then click **Submit**.

 **Note:**

You only need to click Submit once. Use the search field to search for `src_ociggl1` and verify the tables were added.

Task 2: Create the Extracts

Add the Change Data Capture Extract:

In source OCI GoldenGate PostgreSQL deployment details, click Launch Console.

1. From the navigation menu, click Overview.
2. On the Administration Service page, click **Add Extract (plus icon)**.
3. On the Extract Type page, select Change Data Capture Extract, and then click **Next**.
4. Complete the Extract Options as follows, and then click **Next**:
 - a. For **Process Name**, enter a name for the Extract, such as `PSQL`.
 - b. For **Credential Domain**, select Oracle GoldenGate.
 - c. For **Credential Alias**, select the alias.
 - d. For **Begin**, select Now.
 - e. For **Trail Name**, enter a two-character trail name, such as `P1`.
5. On the Extract Parameters page, add the following:

```
TABLE SRC_OCIGGLL.*;
```

6. Click **Create and Run**.

You're returned to the Administration Service Overview page, where you can observe the Extract starting.

Task 3: Create the Distribution Path for Change Data Capture

To create a Distribution Path for Change Data Capture, complete the following:

Create OCI GoldenGate Users and Credentials:

1. In the Oracle Cloud console, on the Deployments page, select the target Big Data deployment.
2. On the deployment details page, click **Launch Console**. Log in with the admin user details created in task 1, step 2.
3. Create a user for the Distribution Path:
 - a. Open the navigation menu, and then click **Administrator**.
 - b. Click **Add New User (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Username**, enter `ggsnet`.
 - For **Role**, select **Operator**.
 - Enter the password twice for verification.
4. In the source PostgreSQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click **Add Credential (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Credential Domain**, enter `GGSNetwork`.
 - For **Credential Alias**, enter `dpuser`.
 - For **Database Name**, you can enter any name and leave the Database Server and Port fields blank, or use the default values.
 - For **User ID**, enter `ggsnet`.
 - For **Password**, enter the same password used in the previous step.
5. Create a Distribution Path:
 - a. In the source MySQL deployment console, click **Distribution Service**, and then click **Add Path (plus icon)**.
 - b. Complete the following fields, and click **Create and Run**:
 - For **Path Name**, enter a name for this path.
 - For **Source Extract**, select the **Change Data Capture Extract (PSQL)**.
 - For **Trail Name**, select the Change Data Capture Extract trail file (P1).
 - For **Target Authentication Method**, select **UserID Alias**.
 - For **Target**, select **wss**.
 - For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the `https://` or any trailing slashes.
 - For **Port Number**, enter **443**.
 - For **Trail Name**, enter `P1`.
 - For **Domain**, enter the domain name created in the previous step.
 - For **Alias**, enter the alias created in the previous step.

You're returned to the Distribution Service Overview page where you can review the path created.

6. In the target Big Data deployment console, click Receiver Service, and then review the Receiver path created.
 - a. Click **Receiver Service**.
 - b. Review the Receiver path details.

Task 4: Add a Replicat for Change Data Capture

Perform updates to the source PostgreSQL database to verify replication to OCI MySQL Database.

1. In the target MySQL deployment console, click **Administrator Service**, and then click **Add Replicat** (plus icon).
2. On the Add Replicat page, under Replicat type, select **Classic**, **Parallel**, or **Coordinated**, and then click **Next**.
3. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - a. For **Process Name**, enter a name, such as `GCPBQ`.
 - b. (Optional) For **Description**, enter a short description to distinguish this process from others.
 - c. For **Trail Name**, enter the name of the Trail from previous task (P1).
 - d. For **Target**, select Google BigQuery from the dropdown.
 - e. For **Available aliases** for Google BigQuery, select your alias from the dropdown.
 - f. For **Available staging locations**, select Google Cloud Storage from the dropdown.
 - g. For **via staging alias**, select Google Cloud Storage connection from the dropdown.
4. On the Parameter Files page, configure the required properties as needed. Look for the ones marked as `#TODO`. And then click **Next**. Some properties to consider modifying include:

```
MAP *.* , TARGET *.*;
```

5. On the Parameter File page, add the following mapping, and then click **Next**:
 - `gg.eventhandler.gcs.bucketMappingTemplate`: provide the name of the bucket that will be used as staging storage
6. Click **Create and Run**.

You return to the Overview page, where you can review the Replicat details.

Task 5: Verify Change Data Capture

Perform updates to the source PostgreSQL database to verify replication to Google BigQuery.

1. Run the following script to perform inserts into the PostgreSQL database:

```
Insert into src_ociggl.src_city  
(CITY_ID,CITY,REGION_ID,POPULATION) values
```

```
(1000, 'Houston', 20, 743113);
Insert into src_ociggl.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1001, 'Dallas', 20, 822416);
Insert into src_ociggl.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1002, 'San Francisco', 21, 157574);
Insert into src_ociggl.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1003, 'Los Angeles', 21, 743878);
Insert into src_ociggl.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1004, 'San Diego', 21, 840689);
Insert into src_ociggl.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1005, 'Chicago', 23, 616472);
Insert into src_ociggl.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1006, 'Memphis', 23, 580075);
Insert into src_ociggl.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1007, 'New York City', 22, 124434);
Insert into src_ociggl.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1008, 'Boston', 22, 275581);
Insert into src_ociggl.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1009, 'Washington D.C.', 22, 688002);
```

2. In the source PostgreSQL deployment console, select the Change Data Capture Extract name (PSQL), and then click **Statistics**. Verify that **src_ociggl.src_city** has 10 inserts.

 **Note:**

If the Extract captured no inserts, then restart the PSQL Extract.

3. In the target Big Data deployment console, select the Change Data Capture Replicat name (GCPBQ), view its **Details**, and check **Statistics** to verify the number of inserts.

SQL Server quickstarts

Common use cases using Microsoft SQL Server databases as OCI GoldenGate sources or targets.

Articles in this section:

- [Replicate data from Azure SQL Managed Instance to Autonomous Transaction Processing](#)

Replicate data from Azure SQL Managed Instance to Autonomous Transaction Processing

Learn to replicate data from Azure SQL Managed Instance to an Autonomous Database using OCI GoldenGate.

Before you begin

To successfully complete this quickstart, you must have the following:

1. [Create an Azure SQL Managed Instance resource.](#)
2. [Configure Azure SQL Managed Instance to allow public connections.](#)

3. Connect to Azure SQL Managed Instance using a SQL client, and then complete the following steps:

- a. Enter the following commands to create a GGADMIN user:

```
CREATE login GGADMIN with password = 'W3lcome@1234'  
Create user GGADMIN for login GGADMIN  
ALTER SERVER ROLE sysadmin ADD MEMBER GGADMIN
```

 **Note:**

The sysadmin role is only required to add trandata and heartbeat tables. You can then remove the sysadmin privileges and instead use db_owner: ALTER ROLE db_owner ADD MEMBER GGADMIN;

- b. Enter the following command to create a database:

```
Create database SRC_OCIGLL
```

- c. Disconnect and reconnect as the newly created user (GGADMIN) and database (SRC_OCIGLL).

- d. To enable Change Data Capture (CDC) at the database level:

```
EXECUTE sys.sp_cdc_enable_db
```

- e. Enter the following command to create the schema:

```
Create schema GGADMIN
```

- f. [Run the SQL script.](#)

4. Set up Autonomous Transaction Processing:

- a. Download and unzip the [sample database schema](#).
- b. In the Oracle Cloud console, select your ATP instance from the Autonomous Databases page to view its details.
- c. Click **Database Actions**.
- d. Enable the GGADMIN user:
 - i. Under **Administration**, click **Database Users**.
 - ii. Locate **GGADMIN**, and then click its ellipsis menu (three dots) and select **Edit**.
 - iii. In the Edit User panel, enter the GGADMIN password, confirm the password, and then deselect **Account is Locked**.
 - iv. Click **Apply Changes**.
- e. Load the target sample schema and data:
 - i. From the Database Actions menu, under **Development**, select **SQL**.
 - ii. Copy and paste the script from `OCIGLL_OCIGGS_SETUP_USERS_ADW.sql` into the SQL worksheet.

- iii. Click **Run Script**. The Script Output tab displays confirmation messages.
- iv. Clear the SQL worksheet and then copy and paste the SQL from `OCIGLL_OCIGGS_SRC_MIRROR_USER_SEED_DATA.sql`.
- v. To verify that the tables were created successfully, close the SQL window and reopen it again. In the Navigator tab, look for the `SRC_OCIGLL` schema and then select tables from their respective dropdowns.

Task 1: Create the OCI GoldenGate Resources

This quickstart example requires deployments and connections for both the source and target.

1. Create deployments for the source SQL Server and target Oracle database.
2. Create an Azure Managed Instance connection for Azure SQL database, using following values:
 - a. For **Database**, enter `SRC_OCIGLL`.
 - b. For **Port**, enter 3342 for public endpoints.
 - c. For **Host**, use the server name from Azure Managed Instance details page in Azure console (for example, `xyz.database.windows.net`).
 - d. For **User**, enter `GGADMIN`.
 - e. For **Password**, enter the `GGADMIN` user password.
 - f. For **Security protocol**, select **Plain** from the dropdown.
3. Assign the Azure Managed Instance connection to the SQL Server deployment.
4. Create a GoldenGate connection and then assign it to the SQL Server deployment.
5. Create an Autonomous Transaction Processing (ATP) connection.
6. Assign the ATP connection to the Oracle deployment.

Task 2: Create the Extracts

1. On the Deployments page, select the SQL Server deployment created in Task 1.
2. On the deployment details page, click **Launch Console**.
3. Sign in to the source SQL Server deployment console using the Administrator credentials specified when you created the deployment in Task 1.
4. Open the navigation menu, and then select **Configuration**.
5. On the **Database** page, in the **Credentials** table, under **Action**, click **Connect to database Azure SQL**. Checkpoint, TRANDATA Information, and Heartbeat options appear.
6. Next to **TRANDATA Information** click **Add TRANDATA**, and complete the following:
 - a. For **Schema Name**, enter `SRC_OCIGLL`.
 - b. Click **Submit**. The deployment console processes your request, but does not refresh the page.

 **Tip:**

The screen will not refresh on submit. To confirm the TRANDATA was added, enter the SRC_OCIGGLL into the search field and then click **Search**. SRC_OCIGGLL is returned and 6 tables are prepared for instantiation.

7. Next, select **Tasks**, select **Purge Change Data** and complete the following:
 - a. Click on the **Add Purge Change Data** (add icon).
 - b. Enter a name.
 - c. Click **Submit**.
8. Next, click **Overview** in the left navigation.
9. On the Overview page, click **Add Extract** (plus icon).
10. On the Add Extract page, complete the following:
 - a. Select **Integrated Load**, and then click **Next**.
 - b. For **Process Name**, enter EINISQL.
 - c. For **Credential Alias**, select a credential from the dropdown menu.
 - d. For **Trail Name**, enter I1.
 - e. Click **Next**.
 - f. On the Parameter File page, in the text area, add a new line to the existing text and add the following: TABLE SRC_OCIGGGLL.*
 - g. Click **Create and Run**. You return to the Overview page.

It may takes a few minutes for the extract to be created. The yellow exclamation point icon changes to a green checkmark.
11. Select the Extract to view its details and review the Report file. It lists all the tables and the number of exported records for each one of them.
12. Click **Overview**.
13. On the Overview page, click **Add Extract** (plus icon).
14. On the Add Extract page, select **Change Data Capture**, and then click **Next**.
15. On the Extract Options page, complete the following:
 - a. For **Process Name**, enter ECDCSQL.
 - b. For **Credential Alias**, select a credential from the dropdown menu.
 - c. For **Trail Name**, enter M1.
 - d. Click **Next**.
16. On the Parameter File page, in the text area, add a new line to the existing text and add the following: TABLE SRC_OCIGGGLL.*
17. Click **Create and Run**. You return to the Overview page.

It may takes a few minutes for the extract to be created. The yellow exclamation point icon changes to a green checkmark.

Task 3: Create the Distribution Paths

1. Create distribution path for Initial Load Extract. In our example, we name our Initial Load Extract, `DPINISQL`.
2. Create distribution path for CDC Extract. In our example, we name our CDC Extract, `DPCDCSQL`.

Task 4: Create the Replicats

1. Launch and log in to the Oracle deployment console created in task 1.
2. Open the navigation menu and then click **Configuration**.
3. On the **Database** page, in the **Credentials** table, under **Action**, click **Connect to Autonomous Transaction Processing**. Checkpoint, TRANDATA Information, and Heartbeat options appear.
4. Next to **TRANDATA Information** click **Add TRANDATA**, and complete the following:
 - a. For **Schema Name**, enter `SRCMIRROR_OCIGLL.CHECKTABLE`.
 - b. Click **Submit**. The deployment console processes your request, but does not refresh the page.

 **Note:**

To verify, click **Search TRANDATA**, and then enter `SRCMIRROR_OCIGLL.CHECKTABLE` into the Search field and click **Search**. `SRCMIRROR_OCIGLL.CHECKTABLE` is returned and 6 tables are prepared for instantiation.

5. Add a checkpoint table for `SRCMIRROR_OCIGLL.CHECKTABLE`.
6. To add a Replicat for Initial Load, complete the following:
 - a. Click **Administrator Service**, and then click **Add Replicat** (plus icon).
 - b. On the Add Replicat page, under Replicat type, select **Nonintegrated Replicat**, and then click **Next**.
 - c. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name, such as `RIL`.
 - For **Credential Domain**, select the domain for the Autonomous Database connection.
 - For **Credential Alias**, select the alias of the Autonomous Database connection.
 - For **Trail Name**, enter the name of the Trail from Task 2 (I1).
 - For **Checkpoint Table**, select the Checkpoint table you created in Step 1.
 - d. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP SRC_OCIGLL.*, TARGET SRCMIRROR_OCIGLL.*;
```

- e. Click **Create and Run**. You're returned to the Overview page, where you can review the Replicat details.
 - f. Select the Initial Load Replicat (**RIL**) and view its **Details**.
 - g. Click **Statistics** and review the number of Inserts. Refresh the page.
 - If the number of Inserts doesn't change, then all the records from the Initial Load have been loaded and you can stop the Replicat (**RIL**).
 - If the number of Inserts continues to increase, then keep refreshing the page until the Initial Load records are all loaded before continuing.
7. To add a Replicat for Change Data Capture, complete the following:
- a. Click **Administrator Service**, and then click **Add Replicat** (plus icon).
 - b. On the Add Replicat page, under Replicat type, select **Nonintegrated Replicat**, and then click **Next**.
 - c. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name, such as RCDC.
 - For **Credential Domain**, select the domain for the Autonomous Database connection.
 - For **Credential Alias**, select the alias of the Autonomous Database connection.
 - For **Trail Name**, enter the name of the Trail from Task 2 (M1).
 - For **Checkpoint Table**, select the Checkpoint table you created in Step 1.
 - d. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP SRC_OCIGLL.*, TARGET SRCMIRROR_OCIGLL.*;
```

- e. Click **Create**. Do not run the Replicat.
- f. On the Administration Service Overview page, select the Replicat for Initial Load (**RIL**) and view its **Details**.
- g. Click **Statistics** and review the number of Inserts. Refresh the page.
 - If the number of Inserts doesn't change, then all the records from the Initial Load have been loaded and you can stop the Replicat (RIL).
 - If the number of Inserts continues to increase, then keep refreshing the page until the Initial Load records are all loaded before continuing.

 **Note:**

If you don't see any Inserts, click **Performance Metrics Service**, select the **Extract**, and then click **Database Statistics**.

- h. Return to the Administration Service Overview page and then start the Replicat for Change Data Capture (RCDC).
- i. After starting the Replicat for Change Data Capture, review its **Details** and **Statistics** to view the number of Inserts.

MySQL quickstarts

Common use cases using MySQL databases as OCI GoldenGate sources or targets.

Articles in this section:

- [Replicate data from OCI MySQL Database to Autonomous Data Warehouse](#)
- [Send Data from MySQL HeatWave to Azure Event Hubs](#)
- [Replicate data from MySQL HeatWave to Amazon Kinesis](#)
- [Replicate Data from MySQL HeatWave to Google Cloud Storage](#)

Replicate data from OCI MySQL Heatwave Database to Autonomous Data Warehouse

Learn how to use OCI GoldenGate to replicate data from OCI MySQL Heatwave Database to Autonomous Data Warehouse.

Overview

This quickstart example demonstrates how to set up and run a replication between OCI MySQL Database and Autonomous Data Warehouse using OCI GoldenGate.

Before you begin

To successfully complete this quickstart, you must have the following:

- OCI Bastion, to connect to OCI MySQL Database, load the sample database, and perform inserts
- An OCI MySQL Database service instance to serve as the source database
- An Autonomous Data Warehouse instance to serve as the target database

Task 0: Set up the environment

To set up the environment for this Quickstart:

1. [Download the sample data script](#), and then run the script on the OCI MySQL database to create the database and load the data.
2. Create a `ggadmin` user using the following script. Remember to replace `<ggadmin-password>` with a valid password:

```
CREATE USER 'ggadmin' IDENTIFIED BY '<ggadmin-password>';  
GRANT SELECT, REPLICATION SLAVE, REPLICATION CLIENT, CREATE,CREATE VIEW,  
EVENT, INSERT, UPDATE, DROP,EXECUTE, DELETE ON *.* TO 'ggadmin';
```

3. Open port 3306, through which OCI GoldenGate can connect.
 - a. In the Oracle Cloud console, locate the subnet that the OCI MySQL database uses.
 - b. In the security list of the subnet, create an Ingress rule for port TCP/3306.

4. [Download the target schema scripts](#), and then run the scripts in the Autonomous Data Warehouse instance to create the schema and tables. You can use the Autonomous Data Warehouse Database Actions SQL tool to run the scripts:
 - a. In the Oracle Cloud console, open the Autonomous Data Warehouse database details page, and then click **Database Actions**.
 - b. In Database Actions, under **Development**, click **SQL**.
 - c. In the SQL tool, copy and paste the script from `OCIGLL_OCIGGS_SETUP_USERS_ADW.sql` into the SQL worksheet, and then click **Run Script**. If successful, the Script Output tab displays confirmation messages.
 - d. Clear the SQL worksheet, and then copy and paste only the Create Table scripts from `OCIGLL_OCIGGS_SRC_MIRROR_USER_SEED_DATA.sql`. The data will be loaded in a later task.

 **Tip:**

You may need to run each Create Table statement separately for the SQL tool to execute the scripts successfully.

- e. To verify the tables were created successfully, close the SQL window and reopen it again. In the Navigator tab, look for the `SRC_OCIGLL` schema, and then select the tables from their respective dropdowns.

Task 1: Create the OCI GoldenGate resources

This quickstart example requires deployments and connections for both the source and target.

1. Create a deployment for the source OCI MySQL database.
2. Create a deployment for the target Autonomous Data Warehouse.
3. Create a connection to the source OCI MySQL database.
4. Create connection for the target Autonomous Data Warehouse.
5. (Optional) If your Oracle deployment doesn't have a public endpoint, then create a connection to GoldenGate, and then assign this connection to the source MySQL deployment.
6. Assign the source connection to the source MySQL deployment.
7. Assign the target connection to the target Oracle deployment.

Task 2: Create the Extracts

1. On the Deployments page, select the MySQL deployment created in Task 1.
2. On the deployment details page, click **Launch Console**.
3. Sign in to the source MySQL deployment console using the Administrator credentials specified when you created the deployment in Task 1.
4. Open the navigation menu, and then select **Configuration**.

5. Review the credentials for the MySQL Database Server connection. Take note of the Domain and Alias.
6. Return to the Administration Service Overview page.
7. On the Administration Service Overview page, click **Add Extract** (plus icon).
8. On the Extract Type page, select **Change Data Capture Extract**, and then click **Next**.
9. Complete the Extract Options as follows, and then click **Next**:
 - For **Process Name**, enter a name for the Extract, such as `CDCEXT`.
 - For **Credential Domain**, select the domain from Step 5.
 - For **Credential Alias**, select the alias from Step 5.
 - For **Trail Name**, enter a two-character trail name, such as `C1`.
10. Enable **Remote** if capturing from a MySQL Database that is not using global transaction identifiers (GTIDs).
11. On the Extract Parameters page, add the following:

```
TABLE SRC_OCIGLL.*;
```

12. Click **Create and Run**.

You're returned to the Administration Service Overview page, where you can observe the Extract starting.

13. Click **Add Extract**.
14. On the Extract Type page, select **Initial Load Extract**, and then click **Next**.
15. Complete the Extract Options as follows, and then click **Next**:
 - For **Process Name**, enter a name, such as `ILEXT`.
 - For **Credential Domain**, select the domain from Step 5
 - For **Credential Alias**, select the alias from Step 5.
 - For **Trail Name**, enter a two-character trail name, such as `I1`.
16. In the Extract Parameters text area, add the following, and then click **Create and Run**:

```
TABLE SRC_OCIGLL.*;
```

Task 3: Create the Distribution Path for Initial Load Extract

1. On the Deployments page, select the target Autonomous Database deployment.
2. On the deployment details page, click **Launch Console**.
3. Create a user for the Distribution Path.
 - a. Open the navigation menu, and then click **Administrator**.
 - b. Click **Add New User** (plus icon), complete the fields as follows, and then click **Submit**:
 - For **Username**, enter `ggsnet`.
 - For **Role**, select **Operator**.

- Enter the password twice for verification.
4. In the source MySQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click **Add Credential** (plus icon), complete the fields as follows, and then click **Submit**:
 - For Credential Domain, enter `GGSNetwork`.
 - For Credential Alias, enter `dpuser`.
 - For Database Name, you can enter any name and leave the Database Server and Port fields blank, or use the default values.
 - For User ID, enter `ggsnet`
 - For Password, enter the same password used in the previous step.
 5. Create a Distribution Path.
 - a. Click **Distribution Service**, and then click **Add Path** (plus icon).
 - b. Complete the following fields, and click **Create and Run**:
 - For **Path Name**, enter a name for this path.
 - For **Source Trail**, leave blank.
 - For **Trail Name**, enter the Initial Load Extract trail name (`I1`).
 - For **Target Authentication Method**, select **UserID Alias**.
 - For **Target**, select **wss**.
 - For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the `https://` or any trailing slashes.
 - For **Port Number**, enter `443`.
 - For **Trail Name**, enter `I1`.
 - For **Domain**, enter the domain name created in the previous step.
 - For **Alias**, enter the alias created in the previous step.

You're returned to the Distribution Service Overview page where you can review the path created.

6. In the target OCI GoldenGate deployment console, review the Receiver path created as a result of the Distribution path:
 - a. Click **Receiver Service**.
 - b. Review the Receiver path details.

Task 4: Add a Replicat for Initial Load

1. In the target OCI GoldenGate deployment console, add a checkpoint table.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click the connect icon for the target Autonomous Database instance.
 - c. Click **Add Checkpoint** (plus icon).

- d. For **Checkpoint Table**, enter `SRCMIRROR_OCIGLL.CHECKTABLE`.
- e. Click **Submit**.

 **Tip:**

The screen will not refresh on submit. To confirm the checkpoint table was added, enter the `SRCMIRROR_OCIGLL.CHECKTABLE` into the search field and then click **Search**.

2. Add the Replicat.
 - a. Click **Administrator Service**, and then click **Add Replicat** (plus icon).
 - b. On the Add Replicat page, under Replicat type, select **Nonintegrated Replicat**, and then click **Next**.
 - c. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name, such as `RIL`.
 - For **Credential Domain**, select the domain for the Autonomous Database connection.
 - For **Credential Alias**, select the alias of the Autonomous Database connection.
 - For **Trail Name**, enter the name of the Trail from Task 2 (I1).
 - For **Checkpoint Table**, select the Checkpoint table you created in Step 1.
 - d. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP SRC_OCIGLL.*, TARGET SRCMIRROR_OCIGLL.*;
```

3. Click **Create and Run**.

You're returned to the Overview page, where you can review the Replicat details.

4. Select the Replicat (**RIL**) and view its **Details**.
5. Click **Statistics** and review the number of Inserts. Refresh the page.
 - If the number of Inserts doesn't change, then all the records from the Initial Load have been loaded and you can stop the Replicat (**RIL**)
 - If the number of Inserts continues to increase, then keep refreshing the page until the Initial Load records are all loaded before continuing.

Task 5: Verify the initial load

1. In the Oracle Cloud console, open **Database Actions** from the Autonomous Data Warehouse database details page.
2. In Database Actions, under **Development**, click **SQL**.
3. In the SQL tool, enter each of the following statements into the worksheet and click **Run Statement**:

```
SELECT * FROM SRCMIRROR_OCIGLL.SRC_CITY;  
SELECT * FROM SRCMIRROR_OCIGLL.SRC_CUSTOMER;
```

The output should return the data that was loaded into the target database tables as a result of the Initial Load.

Task 6: Create a Distribution Path for Change Data Capture

1. In the source MySQL deployment console, click **Distribution Service**.
2. Click **Add Path**.
3. Complete the following fields, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.
 - b. For **Source Extract**, select the Change Data Capture Extract (CDCEXT).
 - c. For **Trail Name**, select the Change Data Capture Extract trail file (C1)
 - d. For **Target Authentication Method**, select **UserID Alias**.
 - e. For **Target**, select **wss**.
 - f. For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the https:// or any trailing slashes.
 - g. For **Port Number**, enter 443.
 - h. For **Trail Name**, enter **C1**.
 - i. For **Domain**, enter the domain name created in task 3.
 - j. For **Alias**, enter the alias created in task 3.
4. In the target OCI GoldenGate deployment console, click **Receiver Service**, and then review the Receiver path created.

Task 7: Add a Replicat for Change Data Capture

1. Add a Replicat.
 - a. Click **Administrator Service**, and then click **Add Replicat** (plus icon).
 - b. On the Add Replicat page, under Replicat type, select **Nonintegrated Replicat**, and then click **Next**.
 - c. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name, such as RCDC.
 - For **Credential Domain**, select the domain for the Autonomous Database connection.
 - For **Credential Alias**, select the alias of the Autonomous Database connection.
 - For **Trail Name**, enter the name of the Trail from Task 2 (C1).
 - For **Checkpoint Table**, select the Checkpoint table you created in Step 1.
 - d. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP SRC_OCIGLL.*, TARGET SRCMIRROR_OCIGLL.*;
```
 - e. Click **Create**. Do not run the Replicat.

2. On the Administration Service Overview page, select the Replicat for Initial Load (**RIL**) and view its **Details**.
3. Click **Statistics** and review the number of Inserts. Refresh the page.
 - If the number of Inserts doesn't change, then all the records from the Initial Load have been loaded and you can stop the Replicat (RIL)
 - If the number of Inserts continues to increase, then keep refreshing the page until the Initial Load records are all loaded before continuing.
4. Return to the Administration Service Overview page and then start the Replicat for Change Data Capture (RCDC).
5. After starting the Replicat for Change Data Capture, review its **Details** and **Statistics** to view the number of Inserts.

Task 8: Verify Change Data Capture

Perform updates to the source OCI MySQL database to verify replication to Autonomous Data Warehouse.

1. In OCI Bastion, create an SSH port forwarding session using MySQL IP and port 3306. Add your public SSH key.
2. Connect to MySQL in Cloud Shell using your private key and port 3306.
3. After connecting successfully, run the following command:

```
mysqlsh admin@localhost:3306 --sql
```

4. Run the following script to perform inserts into the OCI MySQL database:

```
use SRC_OCIGLL;
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1000,'Houston',20,743113);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1001,'Dallas',20,822416);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1002,'San Francisco',21,157574);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1003,'Los Angeles',21,743878);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1004,'San Diego',21,840689);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1005,'Chicago',23,616472);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1006,'Memphis',23,580075);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1007,'New York City',22,124434);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1008,'Boston',22,275581);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1009,'Washington D.C.',22,688002);
commit;
```

5. In the source MySQL deployment console, select the Change Data Capture Extract name (CDCEXT), and then click **Statistics**. Verify that **SRC_OCIGLL.SRC_CITY** has 10 inserts.

 **Note:**

If the Extract captured no inserts, then restart the CDCEXT Extract.

Task 9: Monitor and maintain processes

1. Monitor performance.
2. Manage Trail files.

Send Data from OCI MySQL Heatwave to Azure Event Hubs

Learn how to use OCI GoldenGate to replicate data from OCI MySQL Heatwave to Azure Event Hubs.

Before you begin

To successfully complete this quickstart, you must have the following:

- OCI Bastion, to connect to OCI MySQL Heatwave Database, load the sample database, and perform inserts.
- A OCI MySQL Heatwave Database to serve as the source database.
- [An Azure Event Hubs namespace and an event hub created](#)

 **Note:**

Kafka Surface is not enabled in Azure Event Hubs Basic Tier. Standard or Premium tier is required.

Overview

This quickstart example demonstrates how to set up and run a replication between OCI MySQL Database and Azure Event Hubs using OCI GoldenGate.

Task 0: Set up the environment

To set up the environment for this Quickstart:

1. [Download the sample data script](#), and then run the script on the OCI MySQL Heatwave database to create the database and load the data.
2. Create a `ggadmin` user using the following script. Remember to replace `<ggadmin-password>` with a valid password:

```
CREATE USER 'ggadmin' IDENTIFIED BY '<ggadmin-password>';
GRANT SELECT, REPLICATION SLAVE, REPLICATION CLIENT, CREATE,CREATE
VIEW, EVENT, INSERT, UPDATE, DROP,EXECUTE, DELETE ON *.* TO
'ggadmin';
```

3. Open port 3306, through which OCI GoldenGate can connect.

- a. In the Oracle Cloud console, locate the subnet that the OCI MySQL Heatwave database uses.
- b. In the security list of the subnet, create an Ingress rule for port TCP/3306.

Task 1: Create the OCI GoldenGate resources

1. Create a MySQL deployment for the source OCI MySQL Heatwave database.
2. Create a Big Data deployment for the target Azure Event Hubs.
3. Create a connection to the source MySQL Heatwave Database.
4. Create an Azure Event Hubs connection.
5. Create a connection to GoldenGate, and then assign this connection to the source MySQL deployment.
6. Assign the source connection to the source MySQL deployment.
7. Assign the target connection to the target Big Data deployment.

Task 2: Create the Extracts

1. On the Deployments page, select the MySQL deployment created in Task 1.
2. On the deployment details page, click **Launch Console**.
3. Sign in to the source MySQL deployment console using the Administrator credentials specified when you created the deployment in Task 1.
4. Open the navigation menu, and then select **Configuration**.
5. Review the credentials for the MySQL Database Server connection. Take note of the **Domain** and **Alias**.
6. Return to the Administration Service Overview page.
7. On the Administration Service Overview page, click **Add Extract (plus icon)**.
8. On the Extract Type page, select **Change Data Capture Extract**, and then click **Next**.
9. Complete the Extract Options as follows, and then click **Next**:
 - a. For **Process Name**, enter a name for the Extract, such as CDCEXT.
 - b. For **Credential Domain**, select the domain from Step 5.
 - c. For **Credential Alias**, select the alias from Step 5.
 - d. For **Trail Name**, enter a two-character trail name, such as C1.
10. Enable **Remote** if capturing from a MySQL Database that is not using global transaction identifiers (GTIDs).
11. On the Extract Parameters page, add the following:

```
TABLE SRC_OCIGLL.*;
```

12. Click **Create and Run**.

You're returned to the Administration Service Overview page, where you can observe the Extract starting.

Task 3: Create the Distribution Path to target Big Data deployment

1. On the Deployments page, select the target Big Data deployment.
2. On the deployment details page, click **Launch Console**.
3. Create a user for the Distribution Path:
 - a. Open the navigation menu, and then click **Administrator**.
 - b. Click **Add New User (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Username**, enter `ggsnet`.
 - For **Role**, select **Operator**.
 - Enter the password twice for verification.
4. In the source MySQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click **Add Credential (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Credential Domain**, enter `GGSNetwork`.
 - For **Credential Alias**, enter `dpuser`.
 - For **Database Name**, you can enter any name and leave the Database Server and Port fields blank, or use the default values.
 - For **User ID**, enter `ggsnet`.
 - For **Password**, enter the same password used in the previous step.
5. Create a Distribution Path:
 - a. In the source MySQL deployment console, click **Distribution Service**, and then click **Add Path (plus icon)**.
 - b. Complete the following fields, and click **Create and Run**:
 - For **Path Name**, enter a name for this path.
 - For **Source Extract**, select the **Change Data Capture Extract (CDCEXT)**.
 - For **Trail Name**, select the Change Data Capture Extract trail file (C1).
 - For **Target Authentication Method**, select **UserID Alias**.
 - For **Target**, select **wss**.
 - For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the `https://` or any trailing slashes.
 - For **Port Number**, enter **443**.
 - For **Trail Name**, enter `C1`.
 - For **Domain**, enter the domain name created in the previous step.
 - For **Alias**, enter the alias created in the previous step.

You're returned to the Distribution Service Overview page where you can review the path created.

6. In the target OCI GoldenGate Big Data deployment console, review the Receiver path created as a result of the Distribution path:
 - a. Click **Receiver Service**.
 - b. Review the Receiver path details.

Task 4: Add a Replicat for Azure Event Hubs

1. Add a Replicat:
 - a. In Big Data deployment click **Administrator Service**, and then click **Add Replicat (plus icon)**.
 - b. Select **Classic Replicat** and then click **Next**.
2. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name.
 - b. For **Trail Name**, enter the name of the Trail from Task 2.
 - c. For **Target**, select **Azure Event Hubs**.
 - d. For **Alias**, select the Azure Event Hubs connection created in Task 1.
3. On the Replicat Parameters page, leave the default, and then click **Next**.
4. On the Properties page, provide a topic name for `topicMappingTemplate`.

 **Note:**

The topic name can be a static name or a template keyword for a dynamic topic name. If a topic is not present, OCI GoldenGate generates one for you.

5. Click **Create and Run**.

Task 5: Verify Data Replication from OCI MySQL Heatwave to Azure Event Hubs

Perform updates to the source OCI MySQL Heatwave database to verify replication to Azure Event Hubs.

1. In OCI Bastion, create an SSH port forwarding session using MySQL IP and `port 3306`. Add your public SSH key.
2. Connect to MySQL in Cloud Shell using your private key and `port 3306`.
3. After connecting successfully, run the following command:

```
mysqlsh admin@localhost:3306 --sql
```

4. Run the following script to perform inserts into the OCI MySQL Heatwave database:

```
use SRC_OCIGLL;
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1000,'Houston',20,743113);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1001,'Dallas',20,822416);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
values (1002,'San Francisco',21,157574);
Insert into SRC_OCIGLL.SRC_CITY (CITY_ID,CITY,REGION_ID,POPULATION)
```

```
values (1003,'Los Angeles',21,743878);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1004,'San
Diego',21,840689);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1005,'Chicago',23,616472);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1006,'Memphis',23,580075);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1007,'New York
City',22,124434);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1008,'Boston',22,275581);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1009,'Washington
D.C.',22,688002);
commit;
```

5. In the source MySQL deployment console, select the Change Data Capture Extract name (CDCEXT), and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CITY has 10 inserts.
6. In target Big Data deployment console, select the Replicat Name and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CITY has 10 inserts.
7. In Azure Event Hubs, you can check the message statistics.

Replicate Data from OCI MySQL Heatwave to Amazon Kinesis

Learn how to use OCI GoldenGate to replicate data from OCI MySQL Heatwave to Amazon Kinesis.

Before you begin

To successfully complete this quickstart, you must have the following:

- OCI Bastion, to connect to OCI MySQL Heatwave Database, load the sample database, and perform inserts.
- A OCI MySQL Heatwave Database to serve as the source database.
- [Amazon S3 Access Key & Secret](#)
- [IAM Policies for Amazon Kinesis](#)

Task 0: Set up the environment

To set up the environment for this Quickstart:

1. [Download the sample data script](#), and then run the script on the OCI MySQL Heatwave database to create the database and load the data.

2. Create a `ggadmin` user using the following script. Remember to replace `<ggadmin-password>` with a valid password:

```
CREATE USER 'ggadmin' IDENTIFIED BY '<ggadmin-password>';  
GRANT SELECT, REPLICATION SLAVE, REPLICATION CLIENT, CREATE,CREATE VIEW,  
EVENT, INSERT, UPDATE, DROP,EXECUTE, DELETE ON *.* TO 'ggadmin';
```

3. Open port `3306`, through which OCI GoldenGate can connect.
 - a. In the Oracle Cloud console, locate the subnet that the OCI MySQL Heatwave database uses.
 - b. In the security list of the subnet, create an Ingress rule for port `TCP/3306`.

Task 1: Create the OCI GoldenGate resources

1. Create a MySQL deployment for the source OCI MySQL Heatwave database.
2. Create a Big Data deployment for Amazon Kinesis target.
3. Create a connection to the source MySQL Heatwave Database.
4. Connect to Amazon Kinesis.
5. Create a connection to GoldenGate, and then assign this connection to the source MySQL deployment.
6. Assign the source connection to the source MySQL deployment.
7. Assign the target connection to the target Big Data deployment.

Task 2: Create the Extracts

1. On the Deployments page, select the MySQL deployment created in Task 1.
2. On the deployment details page, click **Launch Console**.
3. Sign in to the source MySQL deployment console using the Administrator credentials specified when you created the deployment in Task 1.
4. Open the navigation menu, and then select **Configuration**.
5. Review the credentials for the MySQL Heatwave connection. Take note of the **Domain** and **Alias**.
6. Return to the Administration Service Overview page.
7. On the Administration Service Overview page, click **Add Extract (plus icon)**.
8. On the Extract Type page, select **Change Data Capture Extract**, and then click **Next**.
9. Complete the Extract Options as follows, and then click **Next**:
 - a. For **Process Name**, enter a name for the Extract, such as `CDCEXT`.
 - b. For **Credential Domain**, select the domain from Step 5.
 - c. For **Credential Alias**, select the alias from Step 5.
 - d. For **Trail Name**, enter a two-character trail name, such as `C1`.
10. Enable **Remote** if capturing from a MySQL Database that is not using global transaction identifiers (GTIDs).

11. On the Extract Parameters page, add the following:

```
TABLE SRC_OCIGLL.*;
```

12. Click **Create and Run**.

You're returned to the Administration Service Overview page, where you can observe the Extract starting.

Task 3: Create the Distribution Path to target Big Data deployment

1. On the Deployments page, select the target Big Data deployment.
2. On the deployment details page, click **Launch Console**.
3. Create a user for the Distribution Path:
 - a. Open the navigation menu, and then click **Administrator**.
 - b. Click **Add New User (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Username**, enter `ggsnet`.
 - For **Role**, select **Operator**.
 - Enter the password twice for verification.
4. In the source MySQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click **Add Credential (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Credential Domain**, enter `GGSTNetwork`.
 - For **Credential Alias**, enter `dpuser`.
 - For **Database Name**, you can enter any name and leave the Database Server and Port fields blank, or use the default values.
 - For **User ID**, enter `ggsnet`.
 - For **Password**, enter the same password used in the previous step.
5. Create a Distribution Path:
 - a. In the source MySQL deployment console, click **Distribution Service**, and then click **Add Path (plus icon)**.
 - b. Complete the following fields, and click **Create and Run**:
 - For **Path Name**, enter a name for this path.
 - For **Source Extract**, select the **Change Data Capture Extract (CDCEXT)**.
 - For **Trail Name**, select the Change Data Capture Extract trail file (C1).
 - For **Target Authentication Method**, select **UserID Alias**.
 - For **Target**, select **wss**.
 - For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the `https://` or any trailing slashes.

- For **Port Number**, enter **443**.
- For **Trail Name**, enter `C1`.
- For **Domain**, enter the domain name created in the previous step.
- For **Alias**, enter the alias created in the previous step.

You're returned to the Distribution Service Overview page where you can review the path created.

6. In the target OCI GoldenGate Big Data deployment console, review the Receiver path created as a result of the Distribution path:
 - a. Click **Receiver Service**.
 - b. Review the Receiver path details.

Task 4: Add a Replicat for Amazon Kinesis

1. Add a Replicat:
 - a. In Big Data deployment click **Administrator Service**, and then click **Add Replicat (plus icon)**.
 - b. Select **Classic Replicat** and then click **Next**.
2. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name.
 - b. For **Trail Name**, enter the name of the Trail from Task 2.
 - c. For **Target**, select **Amazon Kinesis**.
 - d. For **Alias**, select the Amazon Kinesis connection created in Task 1.
3. On the Replicat Parameters page, leave the default, and then click **Next**.
4. On the Properties page, edit the fields marked #TODO.
 - a. `gg.handler.kinesis.region`: provide the Amazon Web Services (AWS) region for the target Kinesis stream.
 - b. `gg.handler.kinesis.streamMappingTemplate`: by default, it is set to `${tableName}` which will map the streams based on source table name. If you want to map to an existing data stream, you can provide static stream names or you can use Template Keywords to assign stream names dynamically.
5. Click **Create and Run**.

Task 5: Verify Data Replication from OCI MySQL Heatwave to Amazon Kinesis

Perform updates to the source OCI MySQL Heatwave database to verify replication to Amazon Kinesis.

1. In OCI Bastion, create an SSH port forwarding session using MySQL IP and `port 3306`. Add your public SSH key.
2. Connect to MySQL in Cloud Shell using your private key and `port 3306`.
3. After connecting successfully, run the following command:

```
mysqlsh admin@localhost:3306 --sql
```

4. Run the following script to perform inserts into the OCI MySQL Heatwave database:

```
use SRC_OCIGLL;
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1000,'Houston',20,743113);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1001,'Dallas',20,822416);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1002,'San
Francisco',21,157574);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1003,'Los
Angeles',21,743878);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1004,'San
Diego',21,840689);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1005,'Chicago',23,616472);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1006,'Memphis',23,580075);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1007,'New York
City',22,124434);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1008,'Boston',22,275581);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1009,'Washington
D.C.',22,688002);
commit;
```

5. In the source MySQL deployment console, select the Change Data Capture Extract name (CDCEXT), and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CITY has 10 inserts.
6. In target Big Data deployment console, select the Replicat Name and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CITY has 10 inserts.
7. In Amazon Kinesis, you can check the messages in target data stream/data viewer.

Replicate Data from OCI MySQL Heatwave to Google Cloud Storage

Learn how to use OCI GoldenGate to replicate data from OCI MySQL Heatwave to Google Cloud Storage.

Before you begin

To successfully complete this quickstart, you must have the following:

- OCI Bastion, to connect to OCI MySQL Heatwave Database, load the sample database, and perform inserts.
- A OCI MySQL Heatwave Database to serve as the source database.
- Google Cloud Storage account and [Google Cloud Service Account Key](#)
- Google Cloud Storage Bucket and Object permissions

Task 0: Set up the environment

To set up the environment for this Quickstart:

1. [Download the sample data script](#), and then run the script on the OCI MySQL Heatwave database to create the database and load the data.
2. Create a `ggadmin` user using the following script. Remember to replace `<ggadmin-password>` with a valid password:

```
CREATE USER 'ggadmin' IDENTIFIED BY '<ggadmin-password>';
GRANT SELECT, REPLICATION SLAVE, REPLICATION CLIENT, CREATE,CREATE VIEW,
EVENT, INSERT, UPDATE, DROP,EXECUTE, DELETE ON *.* TO 'ggadmin';
```

3. Open `port 3306`, through which OCI GoldenGate can connect.
 - a. In the Oracle Cloud console, locate the subnet that the OCI MySQL Heatwave database uses.
 - b. In the security list of the subnet, create an Ingress rule for `port TCP/3306`.

Task 1: Create the OCI GoldenGate resources

1. Create a MySQL deployment for the source OCI MySQL Heatwave database.
2. Create a Big Data deployment for the target Google Cloud Storage.
3. Create a connection to the source MySQL Heatwave Database.
4. Connect to Google Cloud Storage.
5. Create a connection to GoldenGate, and then assign this connection to the source MySQL deployment.
6. Assign the source connection to the source MySQL deployment.
7. Assign the target connection to the target Big Data deployment.

Task 2: Create the Extracts

1. On the Deployments page, select the MySQL deployment created in Task 1.
2. On the deployment details page, click **Launch Console**.
3. Sign in to the source MySQL deployment console using the Administrator credentials specified when you created the deployment in Task 1.
4. Open the navigation menu, and then select **Configuration**.
5. Review the credentials for the MySQL Database Server connection. Take note of the **Domain** and **Alias**.
6. Return to the Administration Service Overview page.
7. On the Administration Service Overview page, click **Add Extract (plus icon)**.

8. On the Extract Type page, select **Change Data Capture Extract**, and then click **Next**.
9. Complete the Extract Options as follows, and then click **Next**:
 - a. For **Process Name**, enter a name for the Extract, such as `CDCEXT`.
 - b. For **Credential Domain**, select the domain from Step 5.
 - c. For **Credential Alias**, select the alias from Step 5.
 - d. For **Trail Name**, enter a two-character trail name, such as `C1`.
10. Enable **Remote** if capturing from a MySQL Database that is not using global transaction identifiers (GTIDs).
11. On the Extract Parameters page, add the following:

```
TABLE SRC_OCIGLL.*;
```

12. Click **Create and Run**.

You're returned to the Administration Service Overview page, where you can observe the Extract starting.

Task 3: Create the Distribution Path to target Big Data deployment

1. On the Deployments page, select the target Big Data deployment.
2. On the deployment details page, click **Launch Console**.
3. Create a user for the Distribution Path:
 - a. Open the navigation menu, and then click **Administrator**.
 - b. Click **Add New User (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Username**, enter `ggsnet`.
 - For **Role**, select **Operator**.
 - Enter the password twice for verification.
4. In the source MySQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click **Add Credential (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Credential Domain**, enter `GGSNetwork`.
 - For **Credential Alias**, enter `dpuser`.
 - For **Database Name**, you can enter any name and leave the Database Server and Port fields blank, or use the default values.
 - For **User ID**, enter `ggsnet`.
 - For **Password**, enter the same password used in the previous step.
5. Create a Distribution Path:
 - a. In the source MySQL deployment console, click **Distribution Service**, and then click **Add Path (plus icon)**.

- b. Complete the following fields, and click **Create and Run**:
 - For **Path Name**, enter a name for this path.
 - For **Source Extract**, select the **Change Data Capture Extract** (CDCEXT).
 - For **Trail Name**, select the Change Data Capture Extract trail file (C1).
 - For **Target Authentication Method**, select **UserID Alias**.
 - For **Target**, select **wss**.
 - For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the https:// or any trailing slashes.
 - For **Port Number**, enter **443**.
 - For **Trail Name**, enter C1.
 - For **Domain**, enter the domain name created in the previous step.
 - For **Alias**, enter the alias created in the previous step.

You're returned to the Distribution Service Overview page where you can review the path created.

6. In the target OCI GoldenGate Big Data deployment console, review the Receiver path created as a result of the Distribution path:
 - a. Click **Receiver Service**.
 - b. Review the Receiver path details.

Task 4: Add a Replicat for Google Cloud Storage

1. Add a Replicat:
 - a. In Big Data deployment click **Administrator Service**, and then click **Add Replicat (plus icon)**.
 - b. Select **Classic Replicat** and then click **Next**.
2. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name.
 - b. For **Trail Name**, enter the name of the Trail from Task 2.
 - c. For **Target**, select **Google Cloud Storage**.
 - d. For **Alias**, select the Azure Event Hubs connection created in Task 1.
3. On the Replicat Parameters page, leave the default, and then click **Next**.
4. On the Properties page, look for `gg.eventhandler.gcs.bucketMappingTemplate=<gcs bucket>`, and replace `<gcs bucket>` with the name of your bucket.

 **Note:**

If target bucket does not exist, OCI GoldenGate creates one for you.

5. Click **Create and Run**.

Task 5: Verify Data Replication from OCI MySQL Heatwave to Google Cloud Storage

Perform updates to the source OCI MySQL Heatwave database to verify replication to Google Cloud Storage.

1. In OCI Bastion, create an SSH port forwarding session using MySQL IP and port 3306. Add your public SSH key.
2. Connect to MySQL in Cloud Shell using your private key and port 3306.
3. After connecting successfully, run the following command:

```
mysqlsh admin@localhost:3306 --sql
```

4. Run the following script to perform inserts into the OCI MySQL Heatwave database:

```
use SRC_OCIGLL;
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1000,'Houston',20,743113);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1001,'Dallas',20,822416);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1002,'San
Francisco',21,157574);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1003,'Los
Angeles',21,743878);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1004,'San
Diego',21,840689);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1005,'Chicago',23,616472);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1006,'Memphis',23,580075);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1007,'New York
City',22,124434);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1008,'Boston',22,275581);
Insert into SRC_OCIGLL.SRC_CITY
(CITY_ID,CITY,REGION_ID,POPULATION) values (1009,'Washington
D.C.',22,688002);
commit;
```

5. In the source MySQL deployment console, select the Change Data Capture Extract name (CDCEXT), and then click **Statistics**. Verify that SRC_OCIGLL.SRC_CITY has 10 inserts.

6. In target Big Data deployment console, select the Replicat Name and then click **Statistics**. Verify that `SRC_OCIGLL.SRC_CITY` has 10 inserts.
7. In Google Cloud Storage, you can check the messages in target data stream/data viewer.

PostgreSQL quickstarts

Common use cases using PostgreSQL databases as OCI GoldenGate sources or targets.

Articles in this section:

- [Replicate Data from PostgreSQL to Autonomous Transaction Processing](#)
- [Replicate data from PostgreSQL to MySQL](#)
- [Replicate data from PostgreSQL to Snowflake](#)
- [Replicate data from PostgreSQL to Google BigQuery](#)

Replicate data from PostgreSQL to Autonomous Transaction Processing

Learn to replicate data from a PostgreSQL server database to Autonomous Transaction Processing using OCI GoldenGate

Before you begin

To successfully complete this quickstart, you must have the following:

- [A PostgreSQL installation](#) to serve as the source database (Installation instructions follow in Task 0.)
- Open port 5432 in your VCN's security list
- An Autonomous Transaction Processing instance to serve as the target database

Task 0: Set up the environment

To set up the environment for this Quickstart:

1. [Install PostgreSQL](#).

- Run the following commands:

- i. Install PostgreSQL server:

```
sudo yum install postgresql-server
```

- ii. Install postgresql-contrib module to avoid [this SQL exception](#):

```
sudo yum install postgresql-contrib
```

- iii. Create a new PostgreSQL database cluster:

```
sudo postgresql-setup --initdb
```

- iv. Enable the postgresql.service:

```
sudo systemctl enable postgresql.service
```

- v. Start the postgresql.service:

```
sudo systemctl start postgresql.service
```

2. By default, PostgreSQL only allows local connections. [Allow remote connectivity to PostgreSQL](#).

- a. In `/var/lib/pgsql/data/postgresql.conf`, prepare the database for replication:

- i. Locate and uncomment `listen_addresses = 'localhost'` and change `localhost` to an asterisk (*):

```
listen_addresses = '*'
```

- ii. Set the following parameters as follows:

- `wal_level = logical`
- `max_replication_slots = 1`
- `max_wal_senders = 1`
- `track_commit_timestamp = on`

 **Note:**

Configure `/var/lib/pgsql/data/pg_hba.conf` to ensure that client authentication is set to allow connections from an Oracle GoldenGate host. For example, add the following:

```
#Allow connections from remote hosts
host    all    all    0.0.0.0/0    md5
```

See [The pg_hba.conf File](#) for more information.

- b. Restart PostgreSQL server:

```
sudo systemctl restart postgresql.service
```

3. If using Oracle Cloud Compute to host PostgreSQL, open port 5432:

```
sudo firewall-cmd --permanent --add-port=5432/tcp
sudo firewall-cmd --reload
sudo firewall-cmd --list-all
```

4. Open port 5432 in your VCN's security list.

5. [Connect to PostgreSQL](#).

```
> sudo su - postgres
> psql
```

6. Set up PostgreSQL.

- a. Download and run [seedSRCOCIGLL_PostgreSQL.sql](#) to set up the database and load the sample data.
- b. Run the following commands to set up the user:

```
create user ggadmin with password 'W3lcome@1234';
alter user ggadmin with SUPERUSER;
GRANT ALL PRIVILEGES ON DATABASE ociggl1 TO ggadmin;
```

7. Set up Autonomous Transaction Processing:
 - a. Download and unzip the [sample database schema](#).
 - b. In the Oracle Cloud console, select your ATP instance from the Autonomous Databases page to view its details and access Database Actions.
 - c. Click **Database Actions**.
 - d. Enable the GGADMIN user:
 - i. Under **Administration**, click **Database Users**.
 - ii. Locate **GGADMIN**, and then click its ellipsis menu (three dots) and select **Edit**.
 - iii. In the Edit User panel, enter the GGADMIN password, confirm the password, and then disable **Account is Locked**.
 - iv. Click **Apply Changes**.
 - e. Load the target sample schema and data:
 - i. From the Database Actions Selector menu, under **Development**, select **SQL**.
 - ii. Copy and paste the script from **OCIGLL_OCIGGS_SETUP_USERS_ADW.sql** into the SQL worksheet.
 - iii. Click **Run Script**. The Script Output tab displays confirmation messages.
 - iv. Clear the SQL worksheet and then copy and paste the SQL from **OCIGLL_OCIGGS_SRC_MIRROR_USER_SEED_DATA.sql**.
 - v. To verify that the tables were created successfully, close the SQL window and reopen it again. In the Navigator tab, look for the SRC_OCIGLL schema and then select tables from their respective dropdowns.

Task 1: Create the OCI GoldenGate resources

This quickstart example requires deployments and connections for both the source and target.

1. Create a deployment for the source PostgreSQL database.
2. Create a deployment for the target Autonomous Transaction Processing instance.
3. Create a connection to the source PostgreSQL database.
 - a. For **Type**, ensure that you select **PostgreSQL Server**.
 - b. For **Database name**, enter `ociggl1`.
 - c. For **Host**, enter the public IP of the Compute instance that PostgreSQL runs on.
 - d. For **Port**, enter `5432`.
 - e. For **Username**, enter `ggadmin`.

- f. For **Password**, enter `w3lcome@1234`.
- g. For **Security Protocol**, select **Plain**.
4. Create connection for the target Autonomous Transaction Processing instance.
5. Create a connection to GoldenGate, and then assign this connection to the source PostgreSQL deployment.
6. Assign the source connection to the source PostgreSQL deployment.
7. Assign the target connection to the target Oracle deployment.

Task 2: Enable supplemental logging

To enable supplemental logging:

1. Launch the PostgreSQL GoldenGate deployment console:
 - a. From the Deployments page, select the PostgreSQL deployment to view its details.
 - b. On the PostgreSQL deployment details page, click **Launch console**.
 - c. On the deployment console sign in page, enter the GoldenGate admin credentials provided in Task 1, step 1.
2. After signing in, open the navigation menu, and then click **Configuration**.
3. Click **Connect**. Checkpoint table and TRANDATA fields appear if the connection is successful.
4. Next to TRANDATA Information, click **Add TRANDATA** (plus icon).
5. For **Table Name**, enter `src_ociggl.*`, and then click **Submit**.

Note:

You only need to click Submit once. Use the search field to search for `src_ociggl` and verify the tables were added.

Task 3: Create the Extracts

1. Add the Change Data Capture Extract:
 - a. On the Administration Service page, click **Add Extract** (plus icon).
 - b. On the Extract Type page, select **Change Data Capture Extract**, and then click **Next**.
 - c. Complete the Extract Options as follows, and then click **Next**:
 - For **Process Name**, enter a name for the Extract, such as `ECDCPSQL`.
 - For **Credential Domain**, select **Oracle GoldenGate**.
 - For **Credential Alias**, select the alias.
 - For **Begin**, select **Now**.
 - For **Trail Name**, enter a two-character trail name, such as `P1`.

- d. On the Extract Parameters page, add the following:

```
TABLE SRC_OCIGLL.*;
```

- e. Click **Create and Run**.
2. Add the Initial Load Extract:
 - a. On the Deployments page, select the PostgreSQL deployment created in Task 1.
 - b. On the deployment details page, click **Launch Console**.
 - c. Sign in to the source PostgreSQL deployment console using the Administrator credentials specified when you created the deployment in Task 1.
 - d. On the Administration Service Overview page, click **Add Extract** (plus icon).
 - e. On the Extract Type page, select **Initial Load Extract**, and then click **Next**.
 - f. Complete the Extract Options as follows, and then click **Next**:
 - For **Process Name**, enter a name, such as `EINIPSQL`.
 - For **Credential Domain**, select **Oracle GoldenGate**.
 - For **Credential Alias**, select the alias.
 - For **Trail Name**, enter a two-character trail name, such as `I1`.
 - g. In the Extract Parameters text area, add the following, and then click **Create and Run**:

```
EXTRACT EINIPSQL
USERIDALIAS PostgreSQL_Compute, DOMAIN OracleGoldenGate
EXTFILE I1, PURGE
TABLE src_ocigll.*;
```

 **Note:**

Ensure that you remove the `SOURCEDB` parameter in front of `USERIDALIAS` before you move on.

- h. Click **Create and Run**.

You return to the Administration Service Overview page, where you can observe the Extract starting.

Task 4: Create the Distribution Path for Initial Load Extract

1. On the Deployments page, select the target Autonomous Database deployment.
2. On the deployment details page, click **Launch Console**.
3. Create a user for the Distribution Path.
 - a. Open the navigation menu, and then click **Administrator**.
 - b. Click **Add New User** (plus icon), complete the fields as follows, and then click **Submit**:
 - For **Username**, enter `ggsnet`.

- For **Role**, select **Operator**.
 - Enter the password twice for verification.
4. In the source MySQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigaton menu, and then select **Configuration**.
 - b. Click **Add Credential** (plus icon), complete the fields as follows, and then click **Submit**:
 - For Credential Domain, enter `GGSNetwork`.
 - For Credential Alias, enter `dpuser`.
 - For Database Name, you can enter any name and leave the Database Server and Port fields blank, or use the default values.
 - For User ID, enter `ggsnet`
 - For Password, enter the same password used in the previous step.
 5. Create a Distribution Path.
 - a. Click **Distribution Service**, and then click **Add Path** (plus icon).
 - b. Complete the following fields, and click **Create and Run**:
 - For **Path Name**, enter a name for this path.
 - For **Source Trail**, leave blank.
 - For **Trail Name**, enter the Initial Load Extract trail name (I1).
 - For **Target Authentication Method**, select **UserID Alias**.
 - For **Target**, select **wss**.
 - For **Target Host**, enter the target Autonomous Database deployment URL, without the `https://` or any trailing slashes.
 - For **Port Number**, enter `443`.
 - For **Trail Name**, enter **I1**.
 - For **Domain**, enter the domain name created in the previous step.
 - For **Alias**, enter the alias created in the previous step.

You're returned to the Distribution Service Overview page where you can review the path created.

6. In the target Autonomous Database deployment console, review the Receiver path created as a result of the Distribution path:
 - a. Click **Receiver Service**.
 - b. Review the Receiver path details.

Task 5: Add a Replicat for Initial Load

1. In the target Autonomous Database deployment, add a checkpoint table.
 - a. Click **Administration Service**, and then open the navigation menu, and select **Configuration**.

- b. Click the connect icon for the target Autonomous Database instance.
 - c. Click **Add Checkpoint** (plus icon).
 - d. For **Checkpoint Table**, enter `SRCMIRROR_OCIGLL.CHECKTABLE`.
 - e. Click **Submit**.
 2. Add the Replicat.
 - a. Open the navigation menu, select **Overview**, and then click **Add Replicat** (plus icon).
 - b. On the Add Replicat page, under Replicat type, select **Nonintegrated Replicat**, and then click **Next**.
 - c. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name, such as `RIL`.
 - For **Credential Domain**, select the domain for the Autonomous Database connection.
 - For **Credential Alias**, select the alias of the Autonomous Database connection.
 - For **Trail Name**, enter the name of the Trail from Task 2 (I1).
 - For **Checkpoint Table**, select the Checkpoint table you created in Step 1.
 - d. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP SRC_OCIGLL.*, TARGET SRCMIRROR_OCIGLL.*;
```

3. Click **Create and Run**.

You return to the Overview page, where you can review the Replicat details.
4. Select the Replicat (**RIL**) and view its **Details**.
5. Click **Statistics** and review the number of Inserts. Refresh the page.
 - If the number of Inserts doesn't change, then all the records from the Initial Load have been loaded and you can stop the Replicat (**RIL**)
 - If the number of Inserts continues to increase, then keep refreshing the page until the Initial Load records are all loaded before continuing.

Task 6: Verify the initial load

1. In the Oracle Cloud console, open **Database actions** from the Autonomous Data Warehouse database details page.
2. In Database Actions, under **Development**, click **SQL**.
3. In the SQL tool, enter each of the following statements into the worksheet and click **Run Statement**:

```
SELECT * FROM SRCMIRROR_OCIGLL.SRC_CITY;  
SELECT * FROM SRCMIRROR_OCIGLL.SRC_CUSTOMER;
```

The output should return the data that was loaded into the target database tables as a result of the Initial Load.

Task 7: Create a Distribution Path for Change Data Capture

1. In the source MySQL deployment console, click **Distribution Service**.
2. Click **Add Path**.
3. Complete the following fields, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.
 - b. For **Source Extract**, select the Change Data Capture Extract (ECDCPSQL).
 - c. For **Trail Name**, select the Change Data Capture Extract trail file (P1)
 - d. For **Target Authentication Method**, select **UserID Alias**.
 - e. For **Target**, select **wss**.
 - f. For **Target Host**, enter the target Autonomous Database deployment console URL, without the https:// or any trailing slashes.
 - g. For **Port Number**, enter 443.
 - h. For **Trail Name**, enter **P1**.
 - i. For **Domain**, enter the domain name created in task 3.
 - j. For **Alias**, enter the alias created in task 3.
4. In the target Autonomous Database deployment console, click **Receiver Service**, and then review the Receiver path created.

Task 8: Add a Replicat for Change Data Capture

1. Add a Replicat.
 - a. Click **Administrator Service**, and then click **Add Replicat** (plus icon).
 - b. On the Add Replicat page, under Replicat type, select **Nonintegrated Replicat**, and then click **Next**.
 - c. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - For **Process Name**, enter a name, such as RCDC.
 - For **Credential Domain**, select the domain for the Autonomous Database connection.
 - For **Credential Alias**, select the alias of the Autonomous Database connection.
 - For **Trail Name**, enter the name of the Trail from Task 2 (P1).
 - For **Checkpoint Table**, select the Checkpoint table you created in Step 1.
 - d. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP SRC_OCIGLL.*, TARGET SRCMIRROR_OCIGLL.*;
```
 - e. Click **Create**. Do not run the Replicat.
2. On the Administration Service Overview page, select the Replicat for Initial Load (**RIL**) and view its **Details**.

3. Click **Statistics** and review the number of Inserts. Refresh the page.
 - If the number of Inserts doesn't change, then all the records from the Initial Load have been loaded and you can stop the Replicat (RIL)
 - If the number of Inserts continues to increase, then keep refreshing the page until the Initial Load records are all loaded before continuing.

 **Note:**

If you don't see any Inserts, click **Performance Metrics Service**, select the **Extract**, and then click **Database Statistics**.

4. Return to the Administration Service Overview page and then start the Replicat for Change Data Capture (RCDC).
5. After starting the Replicat for Change Data Capture, review its **Details** and **Statistics** to view the number of Inserts.

Task 9: Verify Change Data Capture

Perform updates to the source PostgreSQL database to verify replication to Autonomous Transaction Processing.

1. Run the following script to perform inserts into the PostgreSQL database:

```
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1000,'Houston',20,743113);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1001,'Dallas',20,822416);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1002,'San Francisco',21,157574);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1003,'Los Angeles',21,743878);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1004,'San Diego',21,840689);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1005,'Chicago',23,616472);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1006,'Memphis',23,580075);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1007,'New York City',22,124434);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1008,'Boston',22,275581);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1009,'Washington D.C.',22,688002);
```

2. In the source PostgreSQL deployment console, select the Change Data Capture Extract name (ECDCPSQL), and then click **Statistics**. Verify that **src_ociggll.src_city** has 10 inserts.

 **Note:**

If the Extract captured no inserts, then restart the ECDCPSQL Extract.

3. In the target Autonomous Database deployment console, select the Change Data Capture Replicat name (RCDC), view its **Details**, and then check **Statistics**. Verify that SRCMIRROR_OCIGLL.SRC_CITY has 10 inserts.

Task 10: Monitor and maintain processes

1. Monitor performance.
2. Manage Trail files.

Replicate data from PostgreSQL to MySQL

Learn to replicate data from PostgreSQL to MySQL using OCI GoldenGate.

Before you begin

To successfully complete this quickstart, you must have the following:

- [A PostgreSQL installation](#) to serve as the source database (Installation instructions follow in Task 0).
- Open port 5432 in your VCN's security list to be able to access PostgreSQL on its default port.
- An OCI MySQL Database to serve as the target database.

Task 0: Set up the environment

To set up the environment for this Quickstart:

1. [Install PostgreSQL](#).

- Run the following commands:

- i. Install PostgreSQL server:

```
sudo yum install postgresql-server
```

- ii. Install postgresql-contrib module to avoid [this SQL exception](#):

```
sudo yum install postgresql-contrib
```

- iii. Create a new PostgreSQL database cluster:

```
sudo postgresql-setup --initdb
```

- iv. Enable the postgresql.service:

```
sudo systemctl enable postgresql.service
```

- v. Start the postgresql.service:

```
sudo systemctl start postgresql.service
```

2. By default, PostgreSQL only allows local connections. [Allow remote connectivity to PostgreSQL](#).

- In `/var/lib/pgsql/data/postgresql.conf`, prepare the database for replication:
 - i. Locate and uncomment `listen_addresses = 'localhost'` and change `localhost` to an asterisk (*):

```
listen_addresses = '*'
```

- ii. Set the following parameters as follows:

- `wal_level = logical`
- `max_replication_slots = 1`
- `max_wal_senders = 1`
- `track_commit_timestamp = on`

 **Note:**

Configure `/var/lib/pgsql/data/pg_hba.conf` to ensure that client authentication is set to allow connections from an Oracle GoldenGate host. For example, add the following:

```
#Allow connections from remote hosts
host    all    all    0.0.0.0/0    md5
```

See [The pg_hba.conf File](#) for more information.

- Restart PostgreSQL server:

```
sudo systemctl restart postgresql.service
```

3. If using Oracle Cloud Compute to host PostgreSQL, open port 5432:

```
sudo firewall-cmd --permanent --add-port=5432/tcp
sudo firewall-cmd --reload
sudo firewall-cmd --list-all
```

4. Open port 5432 in your VCN's security list.

5. [Connect to PostgreSQL](#).

```
> sudo su - postgres
> psql
```

6. Set up PostgreSQL:

- a. Download and run [seedSRCOCIGLL_PostgreSQL.sql](#) to set up the database and load the sample data.
- b. Run the following commands to set up the user:

```
create user ggadmin with password 'W3lcome@1234';
alter user ggadmin with SUPERUSER;
GRANT ALL PRIVILEGES ON DATABASE ocigll TO ggadmin;
```

7. Set up OCI MySQL Database:
 - a. Open **Port 3306**, through which OCI GoldenGate can connect.
 - i. In the Oracle Cloud console, locate the subnet that the OCI MySQL database uses.
 - ii. In the security list of the subnet, create an Ingress rule for Port TCP/3306.
 - b. Create a `ggadmin` user using the following script. Remember to replace `<ggadmin-password>` with a valid password:

```
CREATE USER 'ggadmin' IDENTIFIED BY '<ggadmin-password>';
GRANT SELECT, REPLICATION SLAVE, REPLICATION CLIENT,
CREATE,CREATE VIEW, EVENT, INSERT, UPDATE, DROP,EXECUTE, DELETE
ON *.* TO 'ggadmin';
```
 - c. Create target tables using [sample schema](#).
 - d. Ensure the tables and user have been successfully created.

Task 1: Create the OCI GoldenGate Resources

This quickstart example requires deployments and connections for both the source and target.

1. Create a deployment for the source PostgreSQL database.
2. Create a MySQL deployment for the target OCI MySQL Database.
3. Create a PostgreSQL connection.
 - a. For **Type**, select **PostgreSQL Server** from the dropdown.
 - b. For **Database name**, enter `ociggl`.
 - c. For **Host**, enter the public IP of the Compute instance that PostgreSQL runs on.
 - d. For **Port**, enter `5432`.
 - e. For **Username**, enter `ggadmin`.
 - f. For **Password**, enter `W3lcome@1234`.
 - g. For **Security Protocol**, select **Plain** from the dropdown.
4. Create an OCI MySQL connection.
5. (Optional) If your OCI MySQL Database doesn't have a public endpoint, then create a connection to GoldenGate, and then assign this connection to the source PostgreSQL deployment.
6. Assign the source connection to the source PostgreSQL deployment.
7. Assign the target connection to the target MySQL deployment.

Task 2: Create the Extracts

1. Enable supplemental logging:
 - a. Launch the PostgreSQL GoldenGate deployment console:
 - i. From the Deployments page, select the PostgreSQL deployment to view its details.

- ii. On the PostgreSQL deployment details page, click **Launch console**.
 - iii. On the deployment console sign in page, enter the GoldenGate admin credentials provided in Task 1, step 1.
 - b. After signing in, open the navigation menu, and then click **Configuration**.
 - c. Click **Connect**. Checkpoint table and TRANDATA fields appear if the connection is successful.
 - d. Next to TRANDATA Information, click **Add TRANDATA** (plus icon).
 - e. For **Table Name**, enter `src_ociggl1.*`, and then click **Submit**.

 **Note:**

You only need to click Submit once. Use the search field to search for `src_ociggl1.*` and verify the tables were added.

2. Add the Change Data Capture Extract:
 - a. In the navigation menu, click **Overview**.
 - b. On the Administration Service page, click **Add Extract** (plus icon).
 - c. On the Extract Type page, select **Change Data Capture Extract**, and then click **Next**.
 - d. Complete the Extract Options as follows, and then click **Next**:
 - i. For **Process Name**, enter a name for the Extract, such as `ECDPCSQL`.
 - ii. For **Credential Domain**, select **Oracle GoldenGate**.
 - iii. For **Credential Alias**, select the alias.
 - iv. For **Begin**, select **Now**.
 - v. For **Trail Name**, enter a two-character trail name, such as `P1`.
 - e. On the Extract Parameters page, add the following:

```
TABLE src_ociggl1.*
```
 - f. Click **Create and Run**.
3. Add the Initial Load Extract:
 - a. On the Administration Service Overview page, click **Add Extract** (plus icon).
 - b. On the Extract Type page, select **Initial Load Extract**, and then click **Next**.
 - c. Complete the Extract Options as follows, and then click **Next**:
 - i. For **Process Name**, enter a name for the Extract, such as `EINIPSQL`.
 - ii. For **Credential Domain**, select **Oracle GoldenGate**.
 - iii. For **Credential Alias**, select the alias.
 - iv. For **Trail Name**, enter a two-character trail name, such as `I1`.

- d. In the Extract Parameters text area, add the following:

```
EXTRACT EINIPSQL
USERIDALIAS PostgreSQL_Compute, DOMAIN OracleGoldenGate
EXTFILE I1, PURGE
TABLE src_ociggll.*;
```

 **Note:**

Ensure that you remove the `SOURCEDB` parameter in front of `USERIDALIAS` before you move on.

- e. Click **Create and Run**.

You return to the Administration Service Overview page, where you can observe the Extract starting.

Task 3: Create the Distribution Path for Initial Load Extract

To create a Distribution Path for Initial Load Extract, complete the following:

1. In the Oracle Cloud console, on the Deployments page, select the target MySQL deployment.
2. On the deployment details page, click **Launch Console**, and log in using the admin detailed created in task 1, step 2.
3. Create a user for the Distribution Path.
 - a. Open the navigation menu, and then click **Administrator**.
 - b. Click **Add New User** (plus icon), complete the fields as follows, and then click **Submit**:
 - i. For **Username**, enter `ggsnet`.
 - ii. For **Role**, select **Operator**.
 - iii. Enter the password twice for verification.
4. In the source PostgreSQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click **Add Credential** (plus icon), complete the fields as follows, and then click **Submit**:
 - i. For **Credential Domain**, enter `GGSNetwork`.
 - ii. For **Credential Alias**, enter `dpuser`.
 - iii. For **Database Name**, you can enter any name and leave the Database Server and Port fields blank, or use the default values.
 - iv. For **User ID**, enter `ggsnet`.
 - v. For **Password**, enter the same password used in the previous step.
5. Create a Distribution Path.
 - a. Click **Distribution Service**, and then click **Add Path** (plus icon).

- b. On the Add Path page, complete the following fields, and click **Create and Run**:
 - i. For **Path Name**, enter a name for this path.
 - ii. For **Source Trail**, leave blank.
 - iii. For **Trail Name**, enter the Initial Load Extract trail name (I1).
 - iv. For **Target Authentication Method**, select **UserID Alias**.
 - v. For **Target**, select **wss**.
 - vi. For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the https:// or any trailing slashes.
 - vii. For **Port Number**, enter 443.
 - viii. For **Trail Name**, enter I1.
 - ix. For **Domain**, enter the domain name created in the previous step.
 - x. For **Alias**, enter the alias created in the previous step.

You return to the Distribution Service Overview page where you can review the path created.

6. In the target MySQL deployment console, review the Receiver path created as a result of the Distribution Path:
 - a. Click **Receiver Service**.
 - b. Review the Receiver Path details.

Task 4: Add the Replicat for Initial Load

1. In the target MySQL deployment console, add a checkpoint table.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click the connect icon for the target OCI MySQL Database.
 - c. Click **Add Checkpoint** (plus icon).
 - d. For **Checkpoint Table**, enter `SRCMIRROR_OCIGLL.CHECKTABLE`.
 - e. Click **Submit**.
2. Add the Initial Load Replicat.
 - a. From the navigation menu, click **Overview**, and then click **Add Replicat** (plus icon).
 - b. On the Add Replicat page, for Replicat type, select either **Classic**, **Parallel**, or **Coordinated**, and then click **Next**.
 - c. On the Replicat Options page, complete the following fields, and then click **Next**:
 - i. For **Process Name**, enter a name, such as RIL.
 - ii. For **Credential Domain**, select the domain for the OCI MySQL Database connection.
 - iii. For **Credential Alias**, select the alias of the OCI MySQL Database connection.
 - iv. For **Trail Name**, enter the name of the Trail from Task 2 (I1).
 - v. For **Checkpoint Table**, select the Checkpoint table you created in Step 1.

- d. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP src_ocigll.*, TARGET SRCMIRROR_OCIGLL.*;
```

3. Click **Create and Run**.

You return to the Overview page, where you can review the Replicat details.

4. Verify the Initial Load:

- a. In Cloud Shell, connect to the VCN and subnet used by your MySQL instance.
- b. Once connected, run the following command :

```
mysqlsh <user>@<MySQL DB Private IP>:3306 --sql
```

- c. Run the following script to verify the data:

```
select * from SRCMIRROR_OCIGLL.SRC_CITY;  
select * from SRCMIRROR_OCIGLL.SRC_CUSTOMER;
```

The output should return the data that was loaded into the target database tables as a result of the Initial Load.

Task 5: Create the Distribution Path for Change Data Capture

To create a Distribution Path for Change Data Capture, complete the following:

1. In the source PostgreSQL deployment console, click **Distribution Service**.
2. Click **Add Path**.
3. On the Add Path page, complete the following fields, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.
 - b. For **Source Extract**, select the Change Data Capture Extract (ECDCPSQL).
 - c. For **Trail Name**, select the Change Data Capture Extract trail file (P1).
 - d. For **Target Authentication Method**, select **UserID Alias**.
 - e. For **Target**, select **wss**.
 - f. For **Target Host**, enter the target MySQL deployment console URL, without the https:// or any trailing slashes.
 - g. For **Port Number**, enter 443.
 - h. For **Trail Name**, enter P1.
 - i. For **Domain**, enter the domain name created in task 3.
 - j. For **Alias**, enter the alias created in task 3.
4. In the target MySQL deployment console, click **Receiver Service**, and then review the Receiver Path created.

Task 6: Add a Replicat for Change Data Capture

Perform updates to the source PostgreSQL database to verify replication to OCI MySQL Database.

1. Add the Change Data Capture Replicat.
 - a. In the target MySQL deployment console, click **Administrator Service**, and then click **Add Replicat** (plus icon).
 - b. On the Add Replicat page, under Replicat type, select **Classic**, **Parallel**, or **Coordinated**, and then click **Next**.
 - c. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - i. For **Process Name**, enter a name, such as `RCDC`.
 - ii. For **Credential Domain**, select the domain for the OCI MySQL Database connection.
 - iii. For **Credential Alias**, select the alias of the OCI MySQL Database connection.
 - iv. For **Trail Name**, enter the name of the Trail from Task 2 (P1).
 - v. For **Checkpoint Table**, select the Checkpoint table you created in Step 1.
 - d. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP src_ociggll.*, TARGET SRCMIRROR_OCIGGLL.*;
```

- e. Click **Create and Run**.

You return to the Overview page, where you can review the Replicat details.

2. Verify the Change Data Capture:
 - a. Perform updates to the source PostgreSQL database to verify replication to OCI MySQL Database. Run the following script to perform inserts into the PostgreSQL database:

```
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1000,'Houston',20,743113);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1001,'Dallas',20,822416);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1002,'San Francisco',21,157574);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1003,'Los Angeles',21,743878);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1004,'San Diego',21,840689);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1005,'Chicago',23,616472);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1006,'Memphis',23,580075);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1007,'New York City',22,124434);
Insert into src_ociggll.src_city (CITY_ID,CITY,REGION_ID,POPULATION)
values (1008,'Boston',22,275581);
```

```
Insert into src_ociggl.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1009,'Washington
D.C.',22,688002);
```

- b. In the source PostgreSQL deployment console, select the Change Data Capture Extract name (ECDCPSQL), click **Details**, and then click **Statistics**. Verify that `src_ociggl.src_city` has 10 inserts.

 **Note:**

If the Extract captured no inserts, then restart the ECDCPSQL Extract.

- c. In the target MySQL deployment, select the Change Data Capture Replicat name (RCDC), review its **Details** and **Statistics** to verify the number of Inserts.

Task 7: Monitor and maintain processes

1. Monitor performance.
2. Manage Trail files.

Replicate data from PostgreSQL to Snowflake

Learn to replicate data from PostgreSQL to Snowflake using OCI GoldenGate.

Before you begin

To successfully complete this quickstart, you must have the following:

- A [PostgreSQL installation](#) to serve as the source database (Installation instructions follow in Task 0).
- Open port 5432 in your VCN's security list.
- A [Snowflake database](#) to serve as the target database.

Set up the environment

To set up the environment for this Quickstart:

1. [Install PostgreSQL](#).

- a. Install PostgreSQL server:

```
sudo yum install postgresql-server
```

- b. Install postgresql-contrib module to avoid [this SQL exception](#):

```
sudo yum install postgresql-contrib
```

- c. Create a new PostgreSQL database cluster:

```
sudo postgresql-setup --initdb
```

- d. Enable the postgresql.service:

```
sudo systemctl enable postgresql.service
```

- e. Start the postgresql.service:

```
sudo systemctl start postgresql.service
```

2. By default, PostgreSQL only allows local connections. [Allow remote connectivity to PostgreSQL](#).

- a. In `/var/lib/pgsql/data/postgresql.conf`, prepare the database for replication:

- i. Locate and uncomment `listen_addresses = 'localhost'` and change localhost to an asterisk (*):

```
listen_addresses = '*'
```

- ii. Set the following parameters as follows:

- `wal_level = logical`
- `max_replication_slots = 1`
- `max_wal_senders = 1`
- `track_commit_timestamp = on`

 **Note:**

Configure `/var/lib/pgsql/data/pg_hba.conf` to ensure that client authentication is set to allow connections from an Oracle GoldenGate host. For example, add the following:

```
#Allow connections from remote hosts
host    all    all    0.0.0.0/0    md5
```

See [The pg_hba.conf File](#) for more information.

- b. Restart PostgreSQL server:

```
sudo systemctl restart postgresql.service
```

3. If using Oracle Cloud Compute to host PostgreSQL, open port 5432:

```
sudo firewall-cmd --permanent --add-port=5432/tcp
sudo firewall-cmd --reload
sudo firewall-cmd --list-all
```

4. Open port 5432 in your VCN's security list.

5. [Connect to PostgreSQL](#).

```
> sudo su - postgres
> psql
```

6. Set up PostgreSQL.
 - a. Download and run [seedSRCOCIGLL_PostgreSQL.sql](#) to set up the database and load the sample data.
 - b. Run the following commands to set up the user:

```
create user ggadmin with password 'W3lcome@1234';  
alter user ggadmin with SUPERUSER;  
GRANT ALL PRIVILEGES ON DATABASE ociggl1 TO ggadmin;
```

7. Set up Snowflake:
 - a. [Create a GoldenGate user in Snowflake](#) with appropriate privileges.
 - b. Create target tables using [sample schema](#).
 - c. Ensure the tables and user have been successfully created.

Task 1: Create the OCI GoldenGate Resources

This quickstart example requires deployments and connections for both the source and target.

1. Create a deployment for the source PostgreSQL database.
2. Create a Big Data deployment for the target Snowflake database.
3. Create a PostgreSQL connection with the following values:
 - a. For **Type**, select **PostgreSQL Server** from the dropdown.
 - b. For **Database name**, enter `ociggl1`.
 - c. For **Host**, enter the public IP of the Compute instance that PostgreSQL runs on.
 - d. For **Port**, enter `5432`.
 - e. For **Username**, enter `ggadmin`.
 - f. For **Password**, enter `w3lcome@1234`.
 - g. For **Security Protocol**, select **Plain** from the dropdown.
4. Create a Snowflake connection with the following values:
 - a. For **Connection URL**, enter `jdbc:snowflake://<account_identifier>.snowflakecomputing.com/?warehouse=<warehouse name>&db=OCIGLL`.

 **Note:**

Ensure you replace `<account_identifier>` and `<warehouse name>` with the appropriate values.

- b. For **Authentication Type**, select **Basic authentication** from the dropdown.
- c. For **Username**, enter a name.
- d. For **Password**, enter a password.

5. (Optional) If your Big Data deployment doesn't have a public endpoint, then create a connection to GoldenGate, and then assign this connection to the source PostgreSQL deployment.
6. Assign the source PostgreSQL connection to the PostgreSQL deployment.
7. Assign the Snowflake connection to the target Big Data deployment.

Task 2: Create the Extracts

1. Enable supplemental logging:
 - a. Launch the PostgreSQL GoldenGate deployment console:
 - i. From the Deployments page, select the PostgreSQL deployment to view its details.
 - ii. On the PostgreSQL deployment details page, click **Launch console**.
 - iii. On the deployment console sign in page, enter the GoldenGate admin credentials provided in Task 1, step 1.
 - b. After signing in, open the navigation menu, and then click **Configuration**.
 - c. Click **Connect**. Checkpoint table and TRANDATA fields appear if the connection is successful.
 - d. Next to TRANDATA Information, click **Add TRANDATA** (plus icon).
 - e. For **Table Name**, enter `src_ociggl1.*;`, and then click **Submit**.

 **Note:**

You only need to click Submit once. Use the search field to search for `src_ociggl1` and verify the tables were added.

2. Add the Change Data Capture Extract:
 - a. From the navigation menu, click **Overview**.
 - b. On the Administration Service page, click **Add Extract** (plus icon).
 - c. On the Extract Type page, select **Change Data Capture Extract**, and then click **Next**.
 - d. Complete the Extract Options as follows, and then click **Next**:
 - i. For **Process Name**, enter a name for the Extract, such as `ECDPCSQL`.
 - ii. For **Credential Domain**, select **Oracle GoldenGate**.
 - iii. For **Credential Alias**, select the alias.
 - iv. For **Begin**, select **Now**.
 - v. For **Trail Name**, enter a two-character trail name, such as `P1`.
 - e. On the Extract Parameters page, add the following:

```
TABLE SRC_OCIGGLL.*;
```
 - f. Click **Create and Run**.
3. Add the Initial Load Extract:

- a. On the Administration Service Overview page, click **Add Extract** (plus icon).
- b. On the Extract Type page, select **Initial Load Extract**, and then click **Next**.
- c. Complete the Extract Options as follows, and then click **Next**:
 - i. For **Process Name**, enter a name for the Extract, such as `EINIPSQL`.
 - ii. For **Credential Domain**, select **Oracle GoldenGate**.
 - iii. For **Credential Alias**, select the alias.
 - iv. For **Trail Name**, enter a two-character trail name, such as `I1`.
- d. In the Extract Parameters text area, add the following:

```
EXTRACT EINIPSQL
USERIDALIAS PostgreSQL_Compute, DOMAIN OracleGoldenGate
EXTFILE I1, PURGE
TABLE src_ociggl1.*;
```

 **Note:**

Ensure that you remove the `SOURCEDB` parameter in front of `USERIDALIAS` before you move on.

- e. Click **Create and Run**.
You return to the Administration Service Overview page, where you can observe the Extract starting.

Task 3: Create the Distribution Path for Initial Load Extract

To create a Distribution Path for Initial Load Extract, complete the following:

1. In the Oracle Cloud console, on the Deployments page, select the target Big Data deployment.
2. On the deployment details page, click **Launch Console**. Log in with the admin user details created in task 1, step 2.
3. Create a user for the Distribution Path.
 - a. Open the navigation menu, and then click **Administrator**.
 - b. Click **Add New User** (plus icon), complete the fields as follows, and then click **Submit**:
 - For **Username**, enter `ggsnet`.
 - For **Role**, select **Operator**.
 - Enter the password twice for verification.
4. In the source PostgreSQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigaton menu, and then select **Configuration**.
 - b. Click **Add Credential** (plus icon), complete the fields as follows, and then click **Submit**:
 - For **Credential Domain**, enter `GGStetwork`.

- For **Credential Alias**, enter `dpuser`.
 - For **Database Name**, you can enter any name and leave the Database Server and Port fields blank, or use the default values.
 - For **User ID**, enter `ggsnet`
 - For **Password**, enter the same password used in the previous step.
5. In the source PostgreSQL deployment console, create a Distribution Path.
 - a. Click **Distribution Service**, and then click **Add Path** (plus icon).
 - b. Complete the following fields, and click **Create and Run**:
 - i. For **Path Name**, enter a name for this path.
 - ii. For **Source Trail**, leave blank.
 - iii. For **Trail Name**, enter the Initial Load Extract trail name (I1).
 - iv. For **Target Authentication Method**, select **UserID Alias**.
 - v. For **Target**, select **wss**.
 - vi. For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the `https://` or any trailing slashes.
 - vii. For **Port Number**, enter `443`.
 - viii. For **Trail Name**, enter `I1`.
 - ix. For **Domain**, enter the domain name created in the previous step.
 - x. For **Alias**, enter the alias created in the previous step.

You return to the Distribution Service Overview page where you can review the path created.
 6. In the target Big Data deployment console, review the Receiver path created as a result of the Distribution path:
 - a. Click **Receiver Service**.
 - b. Review the Receiver path details.

Task 4: Add the Replicat for Initial Load

1. In the target Big Data deployment console, add a checkpoint table.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click the connect icon for the target Snowflake database.
 - c. Click **Add Checkpoint** (plus icon).
 - d. For **Checkpoint Table**, enter `SRCMIRROR_OCIGLL.CHECKTABLE`.
 - e. Click **Submit**.
2. Add the Initial Load Replicat.
 - a. In the navigation menu, click **Overview**, and then click **Add Replicat** (plus icon).
 - b. On the Add Replicat page, under Replicat type, select either **Classic** or **Coordinated**, and then click **Next**.
 - c. On the Replicat Options page, complete the following form fields, and then click **Next**:

- i. For **Process Name**, enter a name, such as RIL.
 - ii. For **Credential Domain**, select the domain for the Snowflake connection.
 - iii. For **Credential Alias**, select the alias of the Snowflake connection.
 - iv. For **Trail Name**, enter the name of the Trail from Task 2 (I1).
 - v. For **Target**, select the target **Snowflake** connection from the dropdown.
 - vi. For **Available aliases**, select an alias from the dropdown, such as Snowflake.
 - vii. (Optional) **Enable external storage** to select an available staging location from the dropdown.
- d. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
INSERTALLRECORDS
MAP src_ociggll.src_city, TARGET SRCMIRROR_OCIGLL.SRC_CITY;
MAP src_ociggll.src_region, TARGET SRCMIRROR_OCIGLL.SRC_REGION;
MAP src_ociggll.src_customer, TARGET
SRCMIRROR_OCIGLL.SRC_CUSTOMER;
MAP src_ociggll.src_orders, TARGET SRCMIRROR_OCIGLL.SRC_ORDERS;
MAP src_ociggll.src_order_lines, TARGET
SRCMIRROR_OCIGLL.SRC_ORDER_LINES;
MAP src_ociggll.src_product, TARGET
SRCMIRROR_OCIGLL.SRC_PRODUCT;
```

- e. On the Properties page, review the properties, and then click **Create and Run**.
You return to the Overview page, where you can review the Replicat details.
3. To verify the Initial Load, connect to Snowflake database and run following queries:

```
select * from SRCMIRROR_OCIGLL.SRC_CITY;
select * from SRCMIRROR_OCIGLL.SRC_CUSTOMER;
```

The output should return the data that was loaded into the target database tables as a result of the Initial Load.

Task 5: Create the Distribution Path for Change Data Capture

To create a Distribution Path for Change Data Capture, complete the following:

1. In the source PostgreSQL deployment console, click **Distribution Service**.
2. On the Paths page, click **Add Path**.
3. On the Add Path page, complete the following fields, and then click **Create and Run**:
 - a. For **Path Name**, enter a name.
 - b. For **Source Extract**, select the Change Data Capture Extract (ECDCPSQL).
 - c. For **Trail Name**, select the Change Data Capture Extract trail file (P1).
 - d. For **Target Authentication Method**, select **UserID Alias**.

- e. For **Target**, select **wss**.
 - f. For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the `https://` or any trailing slashes.
 - g. For **Port Number**, enter 443.
 - h. For **Trail Name**, enter P1.
 - i. For **Domain**, enter the domain name created in task 3.
 - j. For **Alias**, enter the alias created in task 3.
4. In the target Big Data deployment console, click **Receiver Service**, and then review the Receiver path created.

Task 6: Add a Replicat for Change Data Capture

Perform updates to the source PostgreSQL database to verify replication to Snowflake.

1. Add the Change Data Replicat.
 - a. In the target Big Data deployment console, click **Administrator Service**, and then click **Add Replicat** (plus icon).
 - b. On the Add Replicat page, under Replicat type, select either **Classic** or **Coordinated**, and then click **Next**.
 - c. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - i. For **Process Name**, enter a name, such as RCDC.
 - ii. For **Credential Domain**, select the domain for the Snowflake connection.
 - iii. For **Credential Alias**, select the alias of the Snowflake connection.
 - iv. For **Trail Name**, enter the name of the Trail from Task 2 (P1).
 - v. For **Target**, select the target **Snowflake** connection from the dropdown.
 - vi. For **Available aliases**, select an alias from the dropdown, such as Snowflake.
 - vii. (Optional) **Enable external storage** to select an available staging location from the dropdown.
 - d. On the Replicat Parameters page, add the following mapping, and then click **Next**:

```
MAP src_ociggll.src_city, TARGET SRCMIRROR_OCIGLL.SRC_CITY;
MAP src_ociggll.src_region, TARGET SRCMIRROR_OCIGLL.SRC_REGION;
MAP src_ociggll.src_customer, TARGET SRCMIRROR_OCIGLL.SRC_CUSTOMER;
MAP src_ociggll.src_orders, TARGET SRCMIRROR_OCIGLL.SRC_ORDERS;
MAP src_ociggll.src_order_lines, TARGET
SRCMIRROR_OCIGLL.SRC_ORDER_LINES;
MAP src_ociggll.src_product, TARGET SRCMIRROR_OCIGLL.SRC_PRODUCT;
```
 - e. On the Properties page, review the properties, and then click **Create and Run**.
You return to the Overview page, where you can review the Replicat details.
2. Verify the Change Data Capture:

- a. Perform updates to the source PostgreSQL database to verify replication to Snowflake. Run the following script to perform inserts into the PostgreSQL database:

```
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1000,'Houston',20,743113);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1001,'Dallas',20,822416);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1002,'San
Francisco',21,157574);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1003,'Los
Angeles',21,743878);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1004,'San
Diego',21,840689);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1005,'Chicago',23,616472);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1006,'Memphis',23,580075);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1007,'New York
City',22,124434);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1008,'Boston',22,275581);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1009,'Washington
D.C.',22,688002);
```

- b. In the source PostgreSQL deployment console, select the Change Data Capture Extract name (RCDC), and then click **Statistics**. Verify that `src_ociggll.src_city` has 10 inserts.

 **Note:**

If the Extract captured no inserts, then restart the ECDCPSQL Extract.

- c. In the target Big Data deployment console, select the Change Data Capture Replicat name (RCDC), review its **Details** and **Statistics** to verify the number of Inserts.

Task 7: Monitor and maintain processes

1. Monitor performance.
2. Manage Trail files.

Replicate data from PostgreSQL to Google BigQuery

Learn how to use OCI GoldenGate to replicate data from PostgreSQL to Google BigQuery.

Before you begin

To successfully complete this quickstart, you must have the following:

- [A PostgreSQL installation](#) to serve as the source database (Installation instructions follow in Task 0.)
- Open port 5432 in your VCN's security list.
- Create a connection to Google Cloud Storage.

 **Note:**

Please ensure that GCS bucket and the BigQuery dataset exist in the same location/region.

- [Google Cloud Service Account Key](#).
- Google Cloud Platform BigQuery Permissions.

Task 0: Set up the environment

To set up the environment for this Quickstart:

1. [Install PostgreSQL](#).

a. Install PostgreSQL server:

```
sudo yum install postgresql-server
```

b. Install postgresql-contrib module to avoid [this SQL exception](#):

```
sudo yum install postgresql-contrib
```

c. Create a new PostgreSQL database cluster:

```
sudo postgresql-setup --initdb
```

d. Enable the postgresql.service:

```
sudo systemctl enable postgresql.service
```

e. Start the postgresql.service:

```
sudo systemctl start postgresql.service
```

2. By default, PostgreSQL only allows local connections. [Allow remote connectivity to PostgreSQL](#).

a. In `/var/lib/pgsql/data/postgresql.conf`, prepare the database for replication:

- i. Locate and uncomment `listen_addresses = 'localhost'` and change `localhost` to an asterisk (*):

```
listen_addresses = '*'
```

- ii. Set the following parameters as follows:

- `wal_level = logical`
- `max_replication_slots = 1`
- `max_wal_senders = 1`
- `track_commit_timestamp = on`

 **Note:**

Configure `/var/lib/pgsql/data/pg_hba.conf` to ensure that client authentication is set to allow connections from an Oracle GoldenGate host. For example, add the following:

```
#Allow connections from remote hosts
host    all    all    0.0.0.0/0    md5
```

See [The pg_hba.conf File](#) for more information.

- b. Restart PostgreSQL server:

```
sudo systemctl restart postgresql.service
```

3. If using Oracle Cloud Compute to host PostgreSQL, open port 5432:

```
sudo firewall-cmd --permanent --add-port=5432/tcp
sudo firewall-cmd --reload
sudo firewall-cmd --list-all
```

4. Open port 5432 in your VCN's security list.

5. [Connect to PostgreSQL](#).

```
> sudo su - postgres
> psql
```

6. Set up PostgreSQL.

- a. Download and run [seedSRCOCIGLL_PostgreSQL.sql](#) to set up the database and load the sample data.

- b. Run the following commands to set up the user:

```
create user ggadmin with password 'W3lcome@1234';
alter user ggadmin with SUPERUSER;
GRANT ALL PRIVILEGES ON DATABASE ocigll TO ggadmin;
```


Task 1: Create the OCI GoldenGate resources

1. Create a deployment for the source PostgreSQL database.
2. Create a Big Data deployment for the target Google BigQuery.
3. Create a connection to the to the target Google BigQuery.
4. Create a connection to the source PostgreSQL database.
 - a. For **Type**, ensure that you select PostgreSQL Server.
 - b. For **Database name**, enter `ociggl1`.
 - c. For **Host**, enter the public IP of the Compute instance that PostgreSQL runs on.
 - d. For **Port**, enter 5432.
 - e. For **Username**, enter `ggadmin`.
 - f. For **Password**, enter a password.
 - g. For **Security Protocol**, select **Plain**.
5. Create a connection to GoldenGate, and then assign this connection to the source PostgreSQL deployment.
6. Assign the source connection to the source PostgreSQL deployment..
7. Assign the target connection to the target Big Data deployment.
8. Enable supplemental logging:
 - a. Launch the PostgreSQL GoldenGate deployment console:
 - i. From the Deployments page, select the PostgreSQL deployment to view its details.
 - ii. On the PostgreSQL deployment details page, click **Launch console**.
 - iii. On the deployment console sign in page, enter the GoldenGate admin credentials provided in Step 1.
 - b. After signing in, open the navigation menu, and then click **Configuration**.
 - c. Click **Connect**. Checkpoint table and TRANDATA fields appear if the connection is successful.
 - d. Next to TRANDATA Information, click **Add TRANDATA** (plus icon).
 - e. For Table Name, enter `src_ociggl1.*;`, and then click **Submit**.

 **Note:**

You only need to click Submit once. Use the search field to search for `src_ociggl1` and verify the tables were added.

Task 2: Create the Extracts

Add the Change Data Capture Extract:

In source OCI GoldenGate PostgreSQL deployment details, click Launch Console.

1. From the navigation menu, click Overview.
2. On the Administration Service page, click **Add Extract (plus icon)**.
3. On the Extract Type page, select Change Data Capture Extract, and then click **Next**.
4. Complete the Extract Options as follows, and then click **Next**:
 - a. For **Process Name**, enter a name for the Extract, such as `PSQL`.
 - b. For **Credential Domain**, select Oracle GoldenGate.
 - c. For **Credential Alias**, select the alias.
 - d. For **Begin**, select Now.
 - e. For **Trail Name**, enter a two-character trail name, such as `P1`.
5. On the Extract Parameters page, add the following:

```
TABLE SRC_OCIGLL.*;
```

6. Click **Create and Run**.

You're returned to the Administration Service Overview page, where you can observe the Extract starting.

Task 3: Create the Distribution Path for Change Data Capture

To create a Distribution Path for Change Data Capture, complete the following:

Create OCI GoldenGate Users and Credentials:

1. In the Oracle Cloud console, on the Deployments page, select the target Big Data deployment.
2. On the deployment details page, click **Launch Console**. Log in with the admin user details created in task 1, step 2.
3. Create a user for the Distribution Path:
 - a. Open the navigation menu, and then click **Administrator**.
 - b. Click **Add New User (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Username**, enter `ggsnet`.
 - For **Role**, select **Operator**.
 - Enter the password twice for verification.
4. In the source PostgreSQL deployment console, create a credential for the user created in the previous step.
 - a. Open the navigation menu, and then select **Configuration**.
 - b. Click **Add Credential (plus icon)**, complete the fields as follows, and then click **Submit**:
 - For **Credential Domain**, enter `GGNetwork`.
 - For **Credential Alias**, enter `dpuser`.
 - For **Database Name**, you can enter any name and leave the Database Server and Port fields blank, or use the default values.

- For **User ID**, enter `ggsnet`.
 - For **Password**, enter the same password used in the previous step.
5. Create a Distribution Path:
 - a. In the source MySQL deployment console, click **Distribution Service**, and then click **Add Path (plus icon)**.
 - b. Complete the following fields, and click **Create and Run**:
 - For **Path Name**, enter a name for this path.
 - For **Source Extract**, select the **Change Data Capture Extract (PSQL)**.
 - For **Trail Name**, select the Change Data Capture Extract trail file (P1).
 - For **Target Authentication Method**, select **UserID Alias**.
 - For **Target**, select **wss**.
 - For **Target Host**, enter the target OCI GoldenGate deployment console URL, without the `https://` or any trailing slashes.
 - For **Port Number**, enter **443**.
 - For **Trail Name**, enter `P1`.
 - For **Domain**, enter the domain name created in the previous step.
 - For **Alias**, enter the alias created in the previous step.

You're returned to the Distribution Service Overview page where you can review the path created.
 6. In the target Big Data deployment console, click Receiver Service, and then review the Receiver path created.
 - a. Click **Receiver Service**.
 - b. Review the Receiver path details.

Task 4: Add a Replicat for Change Data Capture

Perform updates to the source PostgreSQL database to verify replication to OCI MySQL Database.

1. In the target MySQL deployment console, click **Administrator Service**, and then click **Add Replicat (plus icon)**.
2. On the Add Replicat page, under Replicat type, select **Classic**, **Parallel**, or **Coordinated**, and then click **Next**.
3. On the Replicat Options page, complete the following form fields, and then click **Next**:
 - a. For **Process Name**, enter a name, such as `GCPBQ`.
 - b. (Optional) For **Description**, enter a short description to distinguish this process from others.
 - c. For **Trail Name**, enter the name of the Trail from previous task (P1).
 - d. For **Target**, select Google BigQuery from the dropdown.
 - e. For **Available aliases** for Google BigQuery, select your alias from the dropdown.
 - f. For **Available staging locations**, select Google Cloud Storage from the dropdown.
 - g. For **via staging alias**, select Google Cloud Storage connection from the dropdown.

4. On the Parameter Files page, configure the required properties as needed. Look for the ones marked as #TODO. And then click **Next**. Some properties to consider modifying include:

```
MAP *.* , TARGET *.*;
```

5. On the Parameter File page, add the following mapping, and then click **Next**:
 - `gg.eventhandler.gcs.bucketMappingTemplate`: provide the name of the bucket that will be used as staging storage
6. Click **Create and Run**.

You return to the Overview page, where you can review the Replicat details.

Task 5: Verify Change Data Capture

Perform updates to the source PostgreSQL database to verify replication to Google BigQuery.

1. Run the following script to perform inserts into the PostgreSQL database:

```
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1000,'Houston',20,743113);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1001,'Dallas',20,822416);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1002,'San
Francisco',21,157574);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1003,'Los
Angeles',21,743878);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1004,'San
Diego',21,840689);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1005,'Chicago',23,616472);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1006,'Memphis',23,580075);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1007,'New York
City',22,124434);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values
(1008,'Boston',22,275581);
Insert into src_ociggll.src_city
(CITY_ID,CITY,REGION_ID,POPULATION) values (1009,'Washington
D.C.',22,688002);
```

2. In the source PostgreSQL deployment console, select the Change Data Capture Extract name (PSQL), and then click **Statistics**. Verify that `src_ociggll.src_city` has 10 inserts.

 **Note:**

If the Extract captured no inserts, then restart the `PSQL` Extract.

3. In the target Big Data deployment console, select the Change Data Capture Replicat name (`GCPBQ`), view its **Details**, and check **Statistics** to verify the number of inserts.

Stream Analytics quickstarts

Common use cases using OCI GoldenGate Stream Analytics.

Articles in this section:

- [Build a simple Stream Analytics pipeline](#)
- [Replicate data to Stream analytics](#)

Create a File stream pipeline

Learn to create a simple pipeline in Stream Analytics using a file, such as CSV or JSON, for quick prototyping, testing, or proofs of concept use cases.

 **Note:**

Stream Analytics is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.

Task 1: Create the Stream Analytics deployment

To create a Stream Analytics deployment:

1. In the Console navigation menu, click **Oracle Database**, and then select **GoldenGate**.
2. On the Deployments page, click **Create deployment**.
3. In the Create deployment panel, enter a name and optionally, a description.
4. From the Compartment dropdown, select a compartment in which to create the deployment.
5. For **OCPU count** enter the number of Oracle Compute units (OCPUs) to use.

 **Note:**

One OCPU is equivalent to 16GB of memory. 3 to 4 OCPUs is sufficient for one Stream analytics pipeline. For more information, see OCPU management and billing.

6. (Optional) Select **Auto scaling**.

 **Note:**

Auto scaling enables OCI GoldenGate to scale up to three times the number of OCPUs you specify for OCPU Count, up to 24 OCPUs. For example, if you specify your OCPU Count as 2 and enable Auto Scaling, then your deployment can scale up to 6 OCPUs. If you specify your OCPU Count as 20 and enable Auto Scaling, OCI GoldenGate can only scale up to 24 OCPUs.

7. From the **Subnet in <Compartment>** dropdown, select the subnet to which a private endpoint is created from the OCI GoldenGate service tenancy. This ensures that the deployment is always available over this subnet, as long as the policies for this subnet allow access. The private endpoint is only used to access the deployment console, and doesn't provide access to other resources in the subnet.

To select a subnet in a different compartment, click **Change compartment**.

 **Note:**

You can only select a private subnet when creating a deployment.

8. Select a license type.
9. (Optional) Click **Show advanced options** for network options and to add tags.
 - a. In the **Network** tab,
 - i. Select **Enable GoldenGate console public access** to include a public endpoint in addition to a private endpoint, and allow public access to the deployment console for users. If selected, OCI GoldenGate creates a load balancer in your tenancy to create a public IP. Select a subnet in the same VCN as this deployment in which to create the load balancer.

 **Note:**

The load balancer is a resource that comes with an additional cost. You can manage this resource, but ensure that you don't delete the load balancer while your deployment is still in use.

[Learn more about load balancer pricing.](#)

- ii. Select **Customize endpoint** to provide a private fully qualified domain name (FQDN) prefix that you'll use to access the private service console URL. You can also optionally upload an SSL/TLS certificate (.pem) and its corresponding private key, however, password protected certificates are not supported. It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.
A self-signed certificate is generated for you, if you don't provide one.

 **Note:**

It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.

- b. In the **Maintenance** tab:
 - i. Select **Customize maintenance window** to define the start of the maintenance window to upgrade the deployment.
 - ii. (Optional) For **Major release auto-upgrade period in days**, enter the number of days, between 0 and 365.
 - iii. (Optional) For **Bundle release auto-upgrade period in days**, enter the number of days, between 0 and 180 days.
 - iv. (Optional) For **Security patch auto-upgrade period in days**, enter the number of days, between 0 and 14 days.
 - v. Select **Enable interim release auto-upgrade**, and, optionally, enter the number of days.

 **Note:**

Learn more about scheduling upgrades.

- c. In the **Tags** tab, add tags to help track the resources within your tenancy. Click **+ Additional tag** to add more tags. [Learn more about tagging](#).
10. Click **Next**.
 11. Select **Stream analytics** for deployment type.
 12. The Stream analytics technology type is automatically selected for you.
 13. For **GoldenGate instance name**, enter a name for the stream analytics instance.
 14. For **Administrator username**, enter `osaadmin`.
 15. For **Administrator password**, enter a password, and then confirm that password.
 16. Click **Create**.

Task 2: Create and publish the pipeline

1. Launch the Stream Analytics pipeline.
 - a. From the Stream Analytics deployment details page, click **Launch console**.
 - b. Log in to the Stream Analytics deployment console using the Administrator username and password specified when you created the deployment in Task 1, steps 14 and 15.
2. In the Stream Analytics deployment console, click **Catalog**.
3. Create a File Stream.
4. Create a pipeline using the File stream created in Step 3.

Learn more about the Pipeline Editor. Here are some actions you can perform on your File stream pipeline:

- Learn about Stages and how to add them to your pipeline.
 - Correlate stream and references
 - Apply functions
5. Publish the pipeline.

Replicate data into Stream Analytics

Learn to replicate data from OCI GoldenGate into Stream Analytics.



Note:

Stream Analytics is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.

GoldenGate Stream Analytics began as a Complex Event Processing engine that evolved to run atop of runtime frameworks with Apache Spark and Apache Kafka. Stream Analytics can ingest a stream of data from any source such as a database, GoldenGate, Kafka, JMS, REST, or even a file system file. After the data is ingested, you can run analytics on live data.

Before you begin

To successfully complete this quickstart you need:

- A source Autonomous Database with sample data loaded, and supplemental logging enabled.

Tip:

You can download the [OCI GoldenGate sample data](#), if you need sample data to work with.

- First, edit `SETUP_USERS_ATP.sql` and modify the `SRC_OCIGGLL` user's password to remove the special characters.
- Use the Autonomous Database's Database actions SQL tool to run the two scripts to create the user schema and tables.
- Use the SQL tool to enable supplemental logging.

Follow the steps in [Lab 1, Task 3: Load the ATP schema](#) for more details.

- Unlock the GGADMIN user on the source Autonomous Database instance
 1. On the Autonomous Database Details page, select **Database Users** from the **Database actions** menu.

 **Tip:**

Use the Autonomous Database administrator credentials provided when you created the instance to log in, if prompted.

2. Locate the **GGADMIN** user, and then select **Edit** from its ellipsis (three dots) menu.
3. In the Edit User panel, enter a password, confirm that password, and then deselect **Account is Locked**.
4. Click **Apply Changes**.

Task 1: Create the OCI GoldenGate resources

1. Create the OCI GoldenGate deployment for Data replication.
2. Create a connection for the source database.
3. Assign the connection to the deployment.
4. Create and run an Extract.

Task 2: Create the Stream Analytics resources

1. Create the Stream Analytics deployment.
2. Create a Kafka connection using the Kafka instance's public IP, and select **Plaintext** for Security protocol.
3. Create a GoldenGate connection.
4. Assign the connections to the Stream Analytics deployment.

Task 3: Create and run the pipelines

1. Launch the Stream Analytics deployment console.
2. Review the connections in the Stream Analytics deployment console.
 - a. In the Stream Analytics deployment console, click **Catalog**.
 - b. On the Catalog page, review the list of connections. You should see the GoldenGate connection, the Autonomous Database connection, and the Kafka connection.
3. Start the GoldenGate Big Data cluster:
 - a. In the OCI GoldenGate Stream Analytics deployment console, select **System settings** from the **ossaadmin** user menu.
 - b. In the System Setting dialog, click **Manage Clusters**, and then expand **GGDB Cluster**.
 - c. Click **Start Cluster**. Wait until the cluster status is **Running**, and then close the dialog window.
4. Update the GoldenGate connection credentials:

Although the GoldenGate connection is available in the Stream Analytics deployment console, the GoldenGate credentials don't carry over. Update the password and test the connection.

 - a. Click **Catalog**, and then click the **GoldenGate** connection.

- b. In the **Edit Connection** dialog, click **Next**.
 - c. For **GG Username**, enter `oggadmin`.
 - d. For **GG Password**, click **Change password**, and then enter the password provided when you created the OCI GoldenGate deployment for Data Replication in Task 1.
 - e. Click **Test connection**. If successful, click **Save**.
5. Use the GoldenGate Extract to create and start GoldenGate Change Data. Ensure that you use the Extract details provided in Task 1 on the **GG Change Data Details** page.
6. Update the Autonomous Database user name. Database connections are created with the default user, `ggadmin`. Update the username to `SRC_OCIGLL` (if you used the sample data provided) to access its schema and tables.
- a. Click **Catalog**, and then click the **Autonomous Database connection**.
 - b. In the Edit Connection dialog, click **Next**.
 - c. For **Username**, enter `SRC_OCIGLL`.
 - d. For **Password**, enter the `SRC_OCIGLLpassword` you modified in the Before you begin steps at the start of this quickstart.
 - e. Click **Test connection**. If successful, click **Save**.
7. Use the Autonomous Database lookup tables to create References for Customers and Orders.
8. Use the Kafka connection to create Kafka Streams for Customers and Orders.
9. Use the Autonomous Database SQL tool to perform inserts on the source database.

For example, you can run the following inserts:

```

Insert into SRC_OCIGLL.SRC_ORDERS
(ORDER_ID,STATUS,CUST_ID,ORDER_DATE,CUSTOMER) values
(11,'COM',101,to_date('16-AUG-2023','DD-MON-YYYY'),null);
Insert into SRC_OCIGLL.SRC_ORDERS
(ORDER_ID,STATUS,CUST_ID,ORDER_DATE,CUSTOMER) values
(12,'COM',102,to_date('16-AUG-2023','DD-MON-YYYY'),null);
Insert into SRC_OCIGLL.SRC_ORDERS
(ORDER_ID,STATUS,CUST_ID,ORDER_DATE,CUSTOMER) values
(13,'COM',103,to_date('16-AUG-2023','DD-MON-YYYY'),null);
Insert into SRC_OCIGLL.SRC_ORDERS
(ORDER_ID,STATUS,CUST_ID,ORDER_DATE,CUSTOMER) values
(14,'COM',104,to_date('16-AUG-2023','DD-MON-YYYY'),null);
Insert into SRC_OCIGLL.SRC_ORDERS
(ORDER_ID,STATUS,CUST_ID,ORDER_DATE,CUSTOMER) values
(15,'COM',105,to_date('16-AUG-2023','DD-MON-YYYY'),null);
Insert into SRC_OCIGLL.SRC_ORDERS
(ORDER_ID,STATUS,CUST_ID,ORDER_DATE,CUSTOMER) values
(16,'COM',106,to_date('16-AUG-2023','DD-MON-YYYY'),null);
Insert into SRC_OCIGLL.SRC_ORDERS
(ORDER_ID,STATUS,CUST_ID,ORDER_DATE,CUSTOMER) values
(17,'COM',107,to_date('16-AUG-2023','DD-MON-YYYY'),null);

```

```
Insert into SRC_OCIGLL.SRC_ORDERS
(ORDER_ID,STATUS,CUST_ID,ORDER_DATE,CUSTOMER) values
(18, 'COM',201,to_date('16-AUG-2023', 'DD-MON-YYYY'),null);
Insert into SRC_OCIGLL.SRC_ORDERS
(ORDER_ID,STATUS,CUST_ID,ORDER_DATE,CUSTOMER) values
(19, 'COM',202,to_date('16-AUG-2023', 'DD-MON-YYYY'),null);
```

- 10.** Create a pipeline that uses the Kafka stream created in Step 8.
- 11.** Add a Query stage, and then add a Filter, to return only orders where the CUST_ID of the Orders stream match the CUSTID of the Customers stream.
- 12.** Add target stage.
- 13.** Publish the pipeline.

2

Overview

Learn about Oracle Cloud Infrastructure GoldenGate concepts, connectivity, responsibilities, and what's new.

Articles in this section:

- [About Oracle Cloud Infrastructure GoldenGate](#)
- [OCI GoldenGate connectivity](#)
- [Shared responsibility model](#)
- [What's new in Oracle Cloud Infrastructure GoldenGate](#)

About Oracle Cloud Infrastructure GoldenGate

Learn about data replication concepts to help you get started with Oracle Cloud Infrastructure GoldenGate.

Oracle Cloud Infrastructure GoldenGate is a fully managed, native cloud service that moves data in real-time, at scale. OCI GoldenGate processes data as it moves from one or more data management systems to target databases. You can also design, run, orchestrate, and monitor data replication, transform data, analyze streaming data in real time without having to allocate or manage any compute environments.

Watch a [short overview video](#), or take an [interactive service tour](#) to learn more.



OCI GoldenGate concepts

The following concepts are essential for working with the Oracle Cloud Infrastructure GoldenGate service.

- **Compartment:** Organizes and isolates your cloud resources, such as cloud networks, compute instances, block volumes, or OCI GoldenGate deployments, and database registrations. Only users with permission to a compartment can work with the resources within that compartment. Compartments also serve as a security boundary inside OCI GoldenGate. Only deployments and database registrations within the same compartment can access each other.
- **Connection:** Contains the network connectivity details for a data source or target for OCI GoldenGate. Connections support database and non-database technologies such as Oracle, MySQL, Apache Kafka, OCI Object Storage, OCI Streaming, and Oracle GoldenGate Distribution and Receiver servers.
- **Deployment:** A container for your OCI GoldenGate resources, such as the OCI GoldenGate Deployment Console.
- **Deployment type:** Represents a specific replication scenario.

- **Deployment Backup:** A backup of a deployment's current state, retained for 60 days. It can be used to restore a deployment or create a new deployment with the state of the original deployment at the time the backup was taken.
- **Extract:** A process that runs against the source database and extracts, or captures data.
- **Trail:** A series of files on the source, intermediary, and/or target system where Oracle GoldenGate stores the captured changes to support the continuous extraction and replication of database changes.
- **Replicat:** A process that delivers data to a target database. It reads the trail file on the target database, reconstructs the DML or DDL operations, and applies them to the target database.

For more information about Oracle GoldenGate concepts, see *Components of a Data Replication in Oracle GoldenGate*.

Data Transforms concepts

Whether you're new to Data Transforms or have past experience with Oracle Data Integrator, it's helpful to familiarize yourself with these concepts before starting with OCI GoldenGate Data Transforms. See *Terminology Information in the Using Data Transforms* guide to learn more.

Stream analytics concepts

The following concepts are essential for working with OCI GoldenGate Stream Analytics:

- **Connection:** Stores the connectivity information for a source or target technology.
- **Stream:** A continuous flow of dynamic data.
- **Pipeline:** The workflow data from source to target.
- **Business logic:** Various filters and functions you can apply to a pipeline to obtain the precise data you want to analyze.
- **Publishing:** Makes the pipeline available to all Stream analytics users and sends data to targets.

User roles

Service administrators, Application DBAs, and Data engineers are among the types of data professionals who use OCI GoldenGate. You may perform one or more of the following roles:

- **Service administrators** are responsible for the administration, management, monitoring, diagnostics, lifecycle management, and security for OCI GoldenGate.
- **Application DBAs** use OCI GoldenGate to develop, build, and test solutions focused on high availability, transaction replication, and data warehouse loading.
- **Data engineers** use OCI GoldenGate to develop, build, and test solutions focused on data lake pipelines and stream processing.

OCI GoldenGate Connectivity

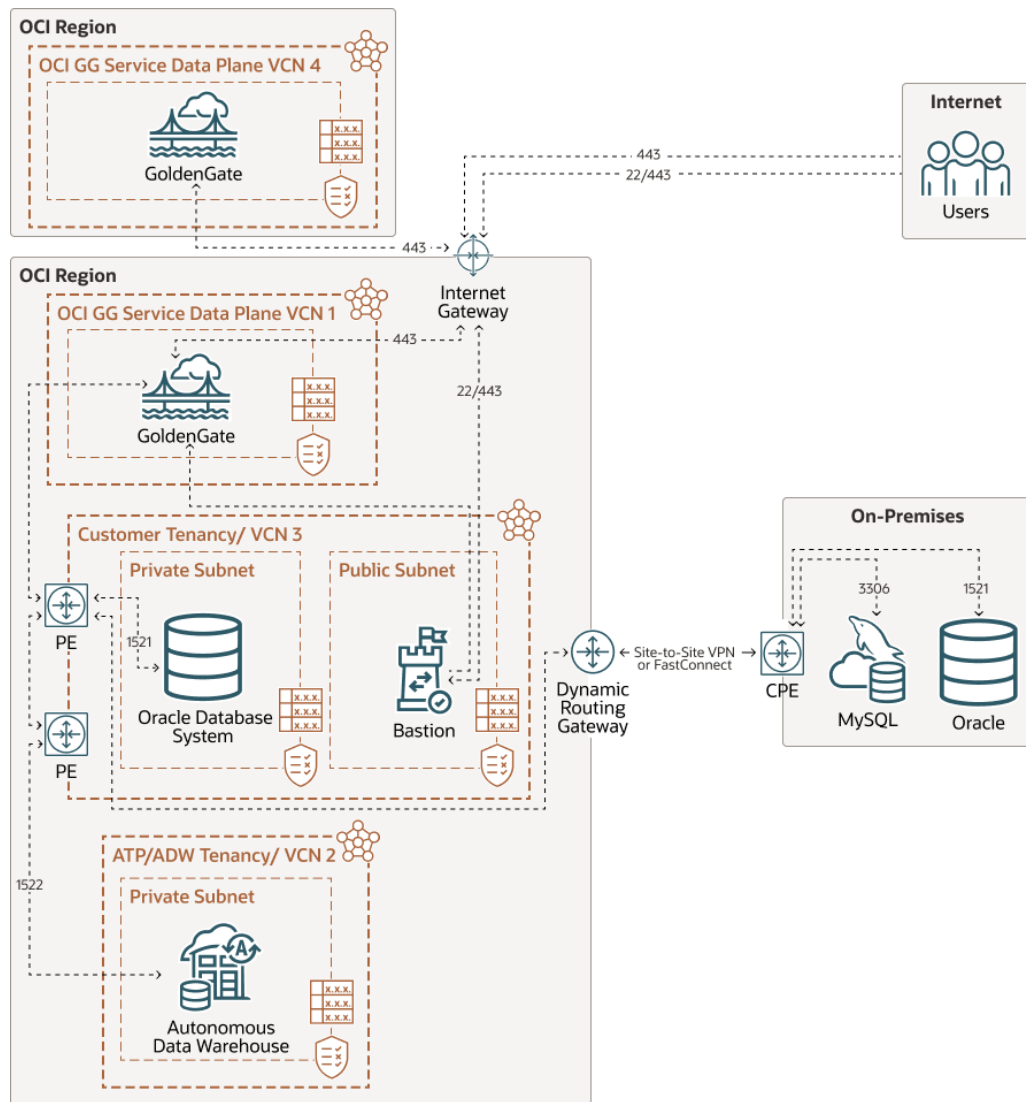
Discover how OCI GoldenGate connects with other services and technologies.

By default, all network connectivity to OCI GoldenGate is encrypted over SSL. The following network diagrams illustrate how traffic is routed through specific ports, depending on the service or technology OCI GoldenGate is connecting to.

Oracle Cloud Infrastructure (OCI) GoldenGate deployments are located in a secure area in OCI outside of your tenancy. Its network isn't connected to any customers' networks and can only access resources available over the public internet by default.

OCI GoldenGate deployments are accessible using a private endpoint and port 443 from a machine using the same subnet. This private endpoint is only used to access the GoldenGate console and doesn't provide access to other resources in that subnet. Optionally, you can 'Enable GoldenGate console public access' and OCI GoldenGate creates a Flexible Load Balancer in the subnet of your choosing, in your tenancy, that connects to the OCI GoldenGate deployment, and creates a public IP. If so, network traffic uses port 443 and the [Internet Gateway](#). You can also add Network Security Groups (NSGs) to the subnet to control traffic.

For example, you can access an OCI GoldenGate deployment through its deployment console. Connectivity to the deployment console is done over HTTPS through port 443. OCI GoldenGate connects to Oracle Databases using the default ports 1521 or 1522, and MySQL databases using default port 3306. For Big Data targets, OCI GoldenGate connects using port 443.



When you create connections, you must specify a Traffic Routing Method. Your choices are:

- **Shared endpoint**, which routes traffic through an endpoint shared with the assigned deployment, and the connection uses the assigned deployment's Ingress IPs and Network Security Group (NSG) settings.

 **Note:**

Shared endpoints may impact performance as connections share bandwidth with the assigned deployment (and potentially other connections assigned to the deployment, if configured to do so).

- **Dedicated endpoint**, which routes traffic through a single dedicated endpoint created in the assigned subnet of your VCN. You must allow connectivity from this connection's Ingress IPs. Select this option to connect to private resources.

 **Note:**

Up to 3 Ingress IPs may be assigned to a connection with a dedicated endpoint. While dedicated endpoint connections don't share bandwidth, they use more IP resources from your resource pool.

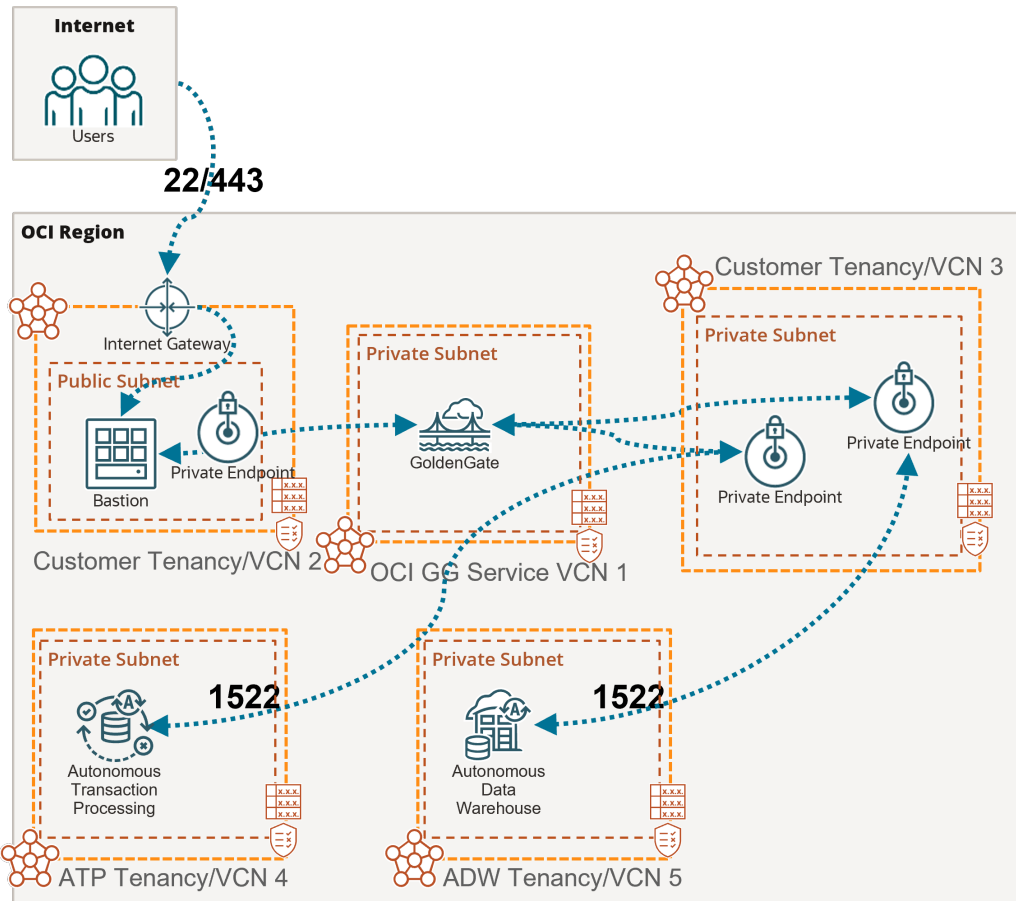
You can connect to publicly available databases with dedicated or shared endpoints, however you must create and configure a Network Access Translation (NAT) Gateway in your VCN. Learn more about [NAT Gateway](#).

When using shared private endpoints, the communication from OCI GoldenGate originates from the Ingress IPs listed on the Connection Details page, after a connection is assigned to a deployment. Ensure that you add the appropriate subnet security rules to allow connectivity from these IP addresses into the data source or target node's private IP.

Any Fully Qualified Domain Name (FQDN) provided must be resolvable within the selected subnet.

Let's look at the following examples.

Example: Replication from Autonomous Transaction Processing into Autonomous Data Warehouse

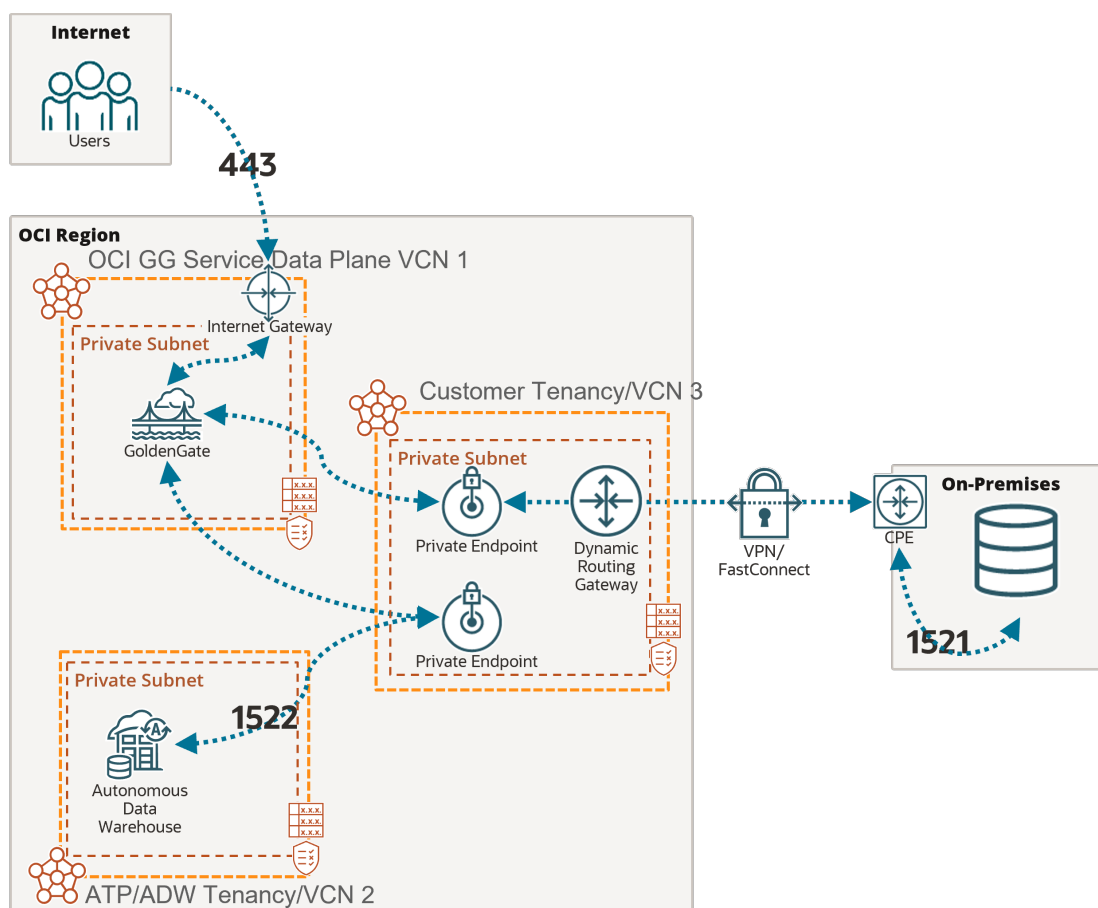


In this example, the OCI GoldenGate deployment is only accessible using a private endpoint from within the OCI network, or through a bastion host that secures access to OCI resources. See [Connect to Oracle Cloud Infrastructure GoldenGate using a private IP](#) for more details.

To connect to Autonomous Transaction Processing (ATP) and Autonomous Data Warehouse (ADW), OCI GoldenGate creates [private endpoints](#) over port 1522, unless you selected 'Secure access from everywhere.'

If you select the Autonomous Database when creating the connection, then the private endpoint gets created automatically. Otherwise, you can enter your Autonomous Database configuration manually, and select 'Shared endpoint' to create a private endpoint in the subnet you select. Appropriate subnet security rules and DNS resolution configuration is your responsibility within the selected subnet.

Example: Replication from Oracle on-premise into Autonomous Data Warehouse

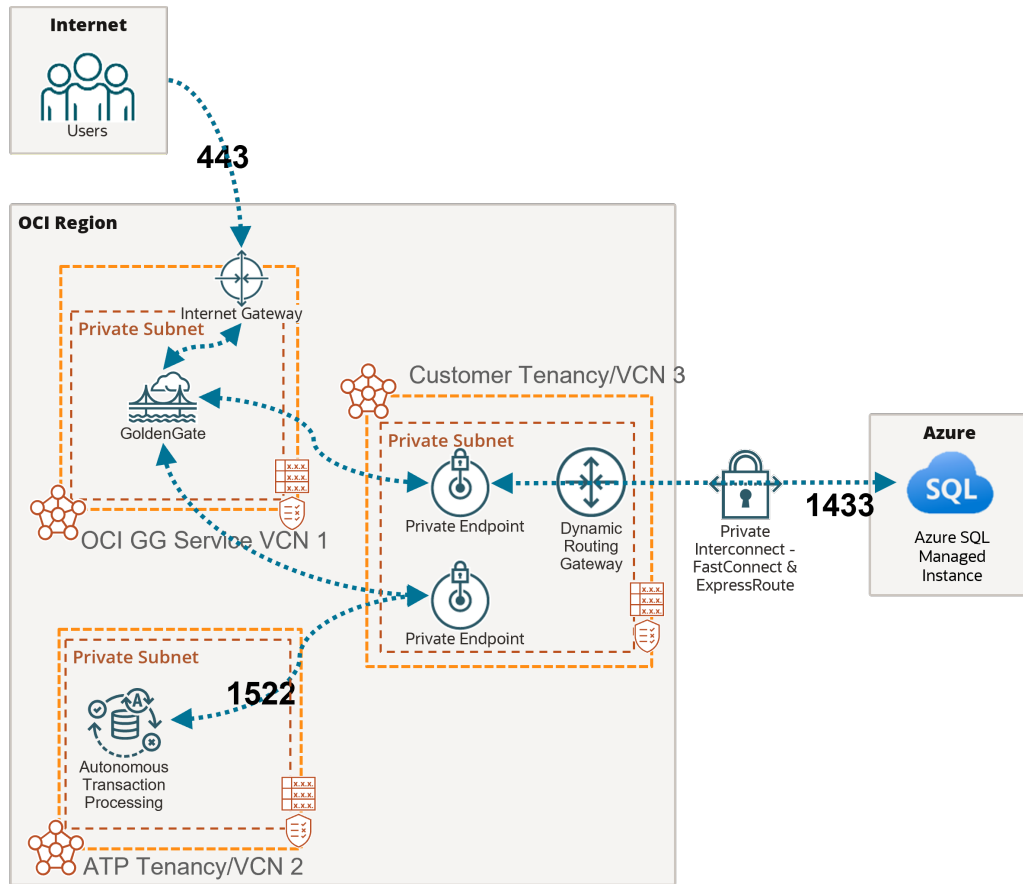


In this example, the OCI GoldenGate deployment is accessible over the public internet using port 443.

To connect to an on-premise Oracle database not available over the public internet, you must create an Oracle connection and select 'Dedicated endpoint' for Traffic routing method. This ensures that OCI GoldenGate creates a [private endpoint](#) in one of your subnets and connects over port 1521. As mentioned above, you must add the appropriate subnet security rules to allow connectivity from the Ingress IPs addresses listed in the Connection Details page into the database node's private IP. Any FQDN provided must be resolvable within the selected subnet.

To connect to Autonomous Data Warehouse (ADW), OCI GoldenGate also creates a [private endpoint](#) over port 1522, unless you selected 'Secure access from everywhere.'

Example: Replication from Azure SQL Managed Instance into Autonomous Transaction Processing



In this example, the OCI GoldenGate deployment is accessible over the public internet using port 443.

To connect to Azure SQL Managed Instance without a public endpoint, a private endpoint must be created. To do so, you must select 'Dedicated endpoint' for Traffic routing method when you create the connection. This ensures that OCI GoldenGate creates a [private endpoint](#) in one of your subnets. As mentioned above, you must add the appropriate subnet security rules to allow connectivity from the Ingress IPs addresses listed on the Connection details page into the database node's private IP. Any FQDN provided must be resolvable within the selected subnet.

In this example, network traffic goes through a private interconnection between OCI and Azure with FastConnect and ExpressRoute. Port 1433 is used for private connections. Connections to Azure SQL Managed Instance over the public internet typically [uses port 3342](#).

To connect to Autonomous Transaction Processing (ATP), OCI GoldenGate creates a [private endpoint](#) over port 1522, unless you selected 'Secure access from everywhere.'

Shared responsibility model

Learn how management tasks for OCI GoldenGate are shared between Oracle and you, the customer.

Task	Who	Details
Provisioning OCI GoldenGate resources	Oracle	Oracle is responsible for provisioning resources. You, the customer, are responsible for initiating provisioning requests that specify configuration characteristics of the resource being provisioned.
Backing up OCI GoldenGate deployments	Oracle	Oracle is responsible for backing up OCI GoldenGate deployments on a daily basis and retaining backups for 60 days.
Restoring an OCI GoldenGate deployment	Oracle	Oracle is responsible for restoring an OCI GoldenGate deployment. You, the customer, are responsible for initiating a restore request that specifies which existing backup to restore to.
Patching and upgrading	Oracle	Oracle periodically releases patches and upgrades for OCI GoldenGate. You, the customer, are responsible for initiating upgrade requests.
Autoscaling	Oracle	Oracle is responsible for scaling OCI GoldenGate per autoscaling configuration settings. You, the customer, are responsible for configuring autoscaling for each OCI GoldenGate deployment.
Scaling	Customer	You, the customer, are responsible for configuring scaling/base OCPU settings for each OCI GoldenGate deployment.
Monitoring service health	Oracle	Oracle is responsible for monitoring the health of OCI GoldenGate resources and for ensuring their availability as per published guidelines.
Monitoring OCI GoldenGate processes' health and performance	Customer	You, the customer, are responsible for monitoring the health and performance of your OCI GoldenGate processes at all levels. This responsibility includes monitoring the performance of the OCI GoldenGate Extract, Replicat, Distribution, and Receiver services and associated processes.
Application security	Customer	You, the customer, are responsible for the security of your applications at all levels. This responsibility includes user access to the OCI GoldenGate resources and network access to these resources. Oracle ensures that data stored in OCI GoldenGate is encrypted and ensures that connections to OCI GoldenGate require SSL encryption.
Alerts and notifications	Oracle / Customer	Oracle is responsible for providing an alert and notification feature for service events. You, the customer are responsible for monitoring any OCI GoldenGate alerts that may be of interest.

What's New in Oracle Cloud Infrastructure GoldenGate

Learn about new features and enhancements recently added to improve your OCI GoldenGate experience.

When new and updated features are available for OCI GoldenGate, instances are upgraded in the data centers where Oracle Cloud services are hosted. Changes to the Oracle Cloud Console take affect automatically, but OCI GoldenGate deployments must be manually upgraded to stay within the support window.

To understand the full scope of each release, ensure that you also review the following resources:

- [OCI GoldenGate Release Notes](#)
- [Known Issues for Oracle Cloud Infrastructure GoldenGate](#)
- [Oracle GoldenGate Release Notes](#)
- [Oracle GoldenGate Bugs Fixed and Enhancements](#)

December 2023


Feature or Change	Description
Networking updates	<p>Updates were made to the following areas:</p> <ul style="list-style-type: none"> • OCI GoldenGate creates a load balancer in your subnet when you create a deployment. This is a resource that you can manage, however, you must refrain from deleting the load balancer while the deployment is in use. • When you create a connection, you must select a routing method. You can choose from: <ul style="list-style-type: none"> – Shared endpoint – Dedicated endpoint <p>See OCI GoldenGate connectivity to learn more.</p>
Contextual notifications	<p>You can now set up contextual notifications on the deployment details page. Contextual notifications let stay informed of specific events that occur with your deployment. Learn more.</p>
Truststore certificates	<p>You can now manage Truststore certificates for other GoldenGate deployments to which your deployment communicates. See Manage Truststore certificates to learn more.</p>
REST APIs	<p>The GoldenGate Service REST API reference was refreshed to reflect the recent feature releases and updates.</p>
New GoldenGate versions available	<p>New GoldenGate versions for Oracle, MySQL, PostgreSQL, and Big Data are now available. Ensure that you upgrade your deployments as soon as you can. See OCI GoldenGate versions.</p>

October 2023

Feature or Change	Description
New connections added	<p>Support for the following connection types is now available in deployment version 841 and above, <i>except</i> for regions where Stream Analytics is in Limited Availability:</p> <ul style="list-style-type: none"> • Google BigQuery • Google Cloud Storage • Redis • Amazon Redshift • Amazon Kinesis • Elasticsearch • Google Cloud SQL for SQL Server • SingleStoreDB • SingleStoreDB Cloud <p>For more information see, What's supported.</p>

Feature or Change	Description
IAM with Identity Domains	<p>GoldenGate now supports Identity Access Management (IAM) with Identity Domains. This feature is available in deployment versions _822 and above. To use IAM with Identity Domains, you must add the necessary policies. For more information, see OCI GoldenGate Policies.</p> <p>Learn more:</p> <ul style="list-style-type: none"> • Create a deployment • Manage deployment users
Documentation updates	<p>New articles added:</p> <ul style="list-style-type: none"> • Quickstarts: <ul style="list-style-type: none"> – Replicate Data from MySQL Heatwave to Google Cloud Storage – Replicate data from MySQL Heatwave to Amazon Kinesis – Replicate data from PostgreSQL to Google BigQuery – Send data from MySQL Heatwave to Azure Event Hubs • Connections: <ul style="list-style-type: none"> – Connect to Google BigQuery – Connect to Google Cloud Storage – Connect to Redis – Connect to Amazon Redshift – Connect to Amazon Kinesis – Connect to Elasticsearch – Connect to Google Cloud SQL for SQL Server – Connect to SingleStoreDB – Connect to SingleStoreDB Cloud • Replicats: <ul style="list-style-type: none"> – Add a Replicat for Google BigQuery – Add a Replicat for Google Cloud Storage – Add a Replicat for Amazon Kinesis – Add a Replicat for Redis – Add a Replicat for Azure Event Hubs
REST APIs	<p>The GoldenGate Service REST API reference was refreshed to reflect the recent feature releases and updates.</p>
New GoldenGate versions available	<p>New GoldenGate versions for Oracle, MySQL, PostgreSQL, and Big Data are now available. Ensure that you upgrade your deployments as soon as you can. See OCI GoldenGate versions.</p>

August 2023

Feature or Change	Description
Stream Analytics	<p>You can now create event stream processing pipelines using OCI GoldenGate Stream Analytics. Pipelines can process input from GoldenGate Extracts, Kafka, or other sources, and join, transform, aggregate, or filter events using patterns and machine learning models. You can then output results to event streams, data stores, or operational dashboards.</p> <div style="border: 1px solid #0070c0; padding: 10px; margin-top: 10px;"> <p> Note:</p> <p>Stream Analytics is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.</p> </div> <p>This feature is available in deployment versions _802 and above. Learn more.</p>
Maintenance	<p>Expanded maintenance features are now available in deployment versions _766 and above. Additions include:</p> <ul style="list-style-type: none"> • More detailed maintenance schedule options when creating and editing deployments • Reschedule upgrades • Cancel upgrades
REST APIs	<p>The GoldenGate Service REST API reference was refreshed to reflect the recent feature releases and updates.</p>
New GoldenGate versions available	<p>New GoldenGate versions for Oracle, MySQL, PostgreSQL, and Big Data are now available. Ensure that you upgrade your deployments as soon as you can. See OCI GoldenGate versions.</p>

July 2023

Feature or Change	Description
Copy deployment backups	<p>You can now copy automatic backups to your Oracle Object Storage bucket. Learn more.</p>
Test connections	<p>You can now test the connectivity to source and target connections assigned to a deployment in the Oracle Cloud console, from either the deployment or connection details page. Learn more.</p>
OCI Logging integration	<p>OCI GoldenGate now supports OCI Logging to produce process and error logs. Learn more.</p>
New GoldenGate versions available	<p>New GoldenGate versions for Oracle, MySQL, PostgreSQL, and Big Data are now available. Ensure that you upgrade your deployments as soon as you can. See OCI GoldenGate versions.</p>

June 2023

Feature or Change	Description
New GoldenGate versions available	New GoldenGate versions for Oracle, MySQL, PostgreSQL, and Big Data are now available. Ensure that you upgrade your deployments as soon as you can. See OCI GoldenGate versions.

March 2023

Feature or Change	Description
New deployment type	Support for Microsoft SQL Server is now available. You can select Microsoft SQL Server as a deployment type, and create the following types of connections: <ul style="list-style-type: none"> Azure SQL Database Azure SQL Managed Instance Microsoft SQL Server Amazon RDS for SQL Server To learn more, see: <ul style="list-style-type: none"> Create a deployment Supported technologies
New Big Data connections	Support for the following Big Data connection types is now available: <ul style="list-style-type: none"> Oracle NoSQL (target only) Oracle Oracle Weblogic JMS Amazon S3 (target only) Hadoop Distributed File System Oracle Autonomous JSON Database (target only) Azure Cosmos DB for MongoDB MongoDB Snowflake (target only) See all supported technologies.
New Quickstarts added	Check out the latest quickstarts for step-by-step instructions on how to get started with our new connection types: <ul style="list-style-type: none"> Replicate data from Azure SQL Managed Instance to Autonomous Transaction Processing Replicate data from PostgreSQL to Snowflake Replicate data from Autonomous Transaction Processing to Amazon S3 Replicate data from MongoDB to Autonomous JSON Database Explore quickstarts.
New maintenance features	You can now schedule upgrades, subscribe to upgrade notifications, snooze upgrade notifications, and rollback upgrades. Learn more. You can also review the different builds available to learn what's included in each release. See OCI GoldenGate versions. You must upgrade to version _703 or above to use the maintenance features.
GoldenGate API updates	To support the features released, additions were made to the GoldenGate API .
New GoldenGate version available	New GoldenGate versions for Oracle, MySQL, PostgreSQL, and Big Data are now available. Ensure that you upgrade your deployments as soon as you can. See OCI GoldenGate versions.

February 2023

Feature or Change	Description
Manage master encryption key wallets	You can now manage master encryption key wallets in OCI GoldenGate. You can use master encryption keys to encrypt trail files distributed to other GoldenGate deployments. Learn more. You must upgrade to version _687 or above to use the master encryption key wallet features.
GoldenGate API updates	To support the features released, additions were made to the GoldenGate API .
New quickstarts added	New quickstarts were added: <ul style="list-style-type: none"> Replicate data from Amazon RDS to OCI Object Storage Stage and merge data into Autonomous Data Warehouse using OCI GoldenGate
New GoldenGate version available	New GoldenGate versions for Oracle, MySQL, PostgreSQL, and Big Data are now available. Ensure that you upgrade your deployments as soon as you can. See OCI GoldenGate versions.

December 2022

Feature or Change	Description
New deployment type added	OCI GoldenGate now supports PostgreSQL as a deployment type. Learn more.
New connection types added	New connection types were added for PostgreSQL and Big Data. See what's supported.
New quickstarts added	New quickstarts were added to showcase the new connection types: <ul style="list-style-type: none"> Replicate data from PostgreSQL to Autonomous Transaction Processing Replicate Data from Autonomous Transaction Processing to Confluent Kafka Replicate data from Autonomous Transaction Processing to Azure Data Lake Storage Replicate Data from Autonomous Transaction Processing to Azure Synapse
Collect diagnostics	You can now collect diagnostics to analyze or share information about your OCI GoldenGate deployment. The information collected can be shared with My Oracle Support, should you encounter any issues. Learn more.
GoldenGate API	To support the features released, additions were made to the Deployments API .
Admin Client	You can now launch the Admin Client directly from the Deployment details page. Learn more about Admin Client.
New GoldenGate version available	A new GoldenGate version for Big Data is now available. Ensure that you upgrade your deployments as soon as you can.

November 2022

Feature or Change	Description
Connections and deployment types	You can use connections and deployment types all regions where GoldenGate is available. Learn more about deployment types and supported connections.
New GoldenGate version	A new GoldenGate version for Oracle is now available. Ensure that you upgrade your deployments as soon as you can.

October 2022

Feature or Change	Description
Connections and deployment types	<p>You can use connections and deployment types all regions where GoldenGate is available <i>except</i>:</p> <ul style="list-style-type: none"> • US West (Phoenix) • US East (Ashburn) • Australia East (Sydney) • Germany Central (Frankfurt) <p>Learn more about deployment types and supported connections.</p>

September 2022

Feature or Change	Description
Connections and deployment types	<p>The following deployment types are now available in the US West (San Jose) and France Central (Paris) regions:</p> <ul style="list-style-type: none"> • Oracle Database • Big Data • MySQL <p>Learn more about these deployment types and support connections.</p>
View Trail files	<p>You can now view Trail files for deployments upgraded to GoldenGate version 21.6 or higher on the Deployment details page, under Resources. Trail files build up over time, which directly impacts the Storage utilization calculation you see under Deployment information. Use this information to manage trail files.</p>

August 2022

Feature or Change	Description
New metrics added	<p>Several new metrics were added to the deployment details page so that you can observe the health of GoldenGate processes within the Oracle Cloud console. Learn more:</p> <ul style="list-style-type: none"> • Metrics • Troubleshoot using the Oracle Cloud console • Monitor performance
Admin client support in Cloud Shell	<p>Cloud Shell now supports GoldenGate Admin client. Admin client enables you to create and manage OCI GoldenGate resources from the command line utility. Learn more.</p>
New GoldenGate version available	<p>A new GoldenGate version is now available. Upgrade your OCI GoldenGate as soon as possible to stay within the current support window. To upgrade, click the Upgrade link on your deployment details page. Learn more.</p>

June 2022

Feature or Change	Description
Connections	A new OCI GoldenGate resource type was introduced and replaces Database Registrations for users in the San Jose region only. Explore connections.
Connections API	A catalog of REST APIs was introduced to support new Connections resource type for users in the San Jose region only.
Supported technologies	You can now use MySQL databases and Big Data technologies with OCI GoldenGate for users in the San Jose region only. Learn what's supported.
New quickstarts	For users in the San Jose region, get started with newly supported databases and technologies with the following quickstarts: <ul style="list-style-type: none"> • Replicate data from MySQL to Autonomous Data Warehouse • Replicate data from Autonomous Transaction Processing to Oracle Object Storage • Replicate data from Autonomous Transaction Processing to OCI Streaming • Replicate data from Autonomous Transaction Processing to Apache Kafka
New troubleshooting topic added	Learn to troubleshoot memory consumption issues and configure CacheMgr.

March 2022

Feature or Change	Description
Storage utilization	Each deployment has a soft limit of 250 GB storage per OCPU. You can see how much storage space your deployment uses on the deployment details page. When the storage limit is reached, a message displays on the deployment details page advising you to take action to free up space. Learn to purge unused trail files.
GoldenGate API	To support the features released, additions were made to the Deployments API .
GoldenGate Events	A new Deployment event type was added for Storage Utilization. Learn more.
New quickstart	Learn to configure bidirectional replication.

December 2021

Feature or Change	Description
Support for RAC databases	Database Registrations now support RAC databases. See Registering a Database .
GoldenGate API	To support the features released, additions were made to the DatabaseRegistration APIs.

November 2021

Feature or Change	Description
New details for Deployment Backups	Deployment backup size and start and end times were added to the Deployment Backup Details page. See Viewing Backup Details .
GoldenGate API	To support the features released, additions were made to the DeploymentBackup APIs.

Feature or Change	Description
Policy examples updated	With the recent release of Oracle Cloud Infrastructure IAM and the introduction of identity domains, updates were made to OCI GoldenGate Policy examples. See: <ul style="list-style-type: none"> • Creating Policies • Policy Examples
Quickstarts updated	Send Data from an On-premises Oracle GoldenGate to OCI GoldenGate and Send Data from OCI GoldenGate to an On-premises Oracle GoldenGate quickstarts were updated to reflect changes to the root certificate.
Managing Deployment Backups	Reasons that a backup process fails were added to Managing Deployment Backups.

October 2021

Feature or Change	Description
Upgrade History	You can now view a list of past upgrades on the Deployment Details page. The Upgrade History displays the Oracle GoldenGate version applied, the date and time the upgrade started and finished, and the completion status. See Upgrading Oracle GoldenGate .
DeploymentUpgrade API	A set of API operations was added to support the Upgrade History feature. See DeploymentUpgrade Reference .
Root certificate change	For OCI GoldenGate users who send data between OCI GoldenGate and on-premises Oracle GoldenGate instances or other non-OCI GoldenGate deployments, OCI GoldenGate's root certificate was recently updated and directly impacts any distribution and receiver paths between these source and targets. After upgrading your deployment, complete the following steps to successfully restart any distribution and receiver paths between OCI GoldenGate and on-premises Oracle GoldenGate instances or other non-OCI GoldenGate deployments: <ul style="list-style-type: none"> • For Oracle GoldenGate 19c users, add the new certificate to the Distribution or Receiver server's client wallet. For more information, see Creating a Distribution Server User Certificate. • For Oracle GoldenGate 21c users, use the Service Manager's Certificate Management feature. For information, see Create a Trusted Connection Between Oracle GoldenGate and OCI GoldenGate.
Managing Trail Files	Learn about Trail file management in Managing Trail Files .
New quickstarts	The following quickstarts were added: <ul style="list-style-type: none"> • Securing a Public Deployment • Replicating Data Across Different Regions with VCN Peering
Quickstart updated	Replicating Data Between Two Cloud Databases now includes steps to instantiate a target database using Oracle Data Pump.
New Policy examples	Policy examples for securing network resources were added. See Policy Examples for Securing Network Resources .

3

Plan

Learn the details service administrators must plan for, such as networking, source and target technologies supported, metering and billing, and integrated services, before you can use Oracle Cloud Infrastructure GoldenGate.

Articles in this section:

- [Before you begin with Oracle Cloud Infrastructure GoldenGate](#)
- [Supported technologies](#)
- [OCPU management and billing](#)
- [Integrated services](#)
- [Example OCI GoldenGate topologies](#)

Before you begin with Oracle Cloud Infrastructure GoldenGate

Familiarize yourself with Oracle Cloud concepts, networking, service availability, and limits before you start using OCI GoldenGate.

Learn Oracle Cloud

If you're completely new to Oracle Cloud, then you should review [Getting Started with Oracle Cloud](#) to learn about types of accounts, terminology, and how to sign up.

You should also familiarize yourself with Oracle Cloud Infrastructure concepts such as tenancies, compartments, VCNs and subnets, and policies, before you create any OCI GoldenGate resources. See:

- [Setting up your tenancy](#)
- [Managing compartments](#)
- [VCNs and subnets](#)
- [How policies work](#)

Region availability

See [OCI Data Regions](#) to discover in which regions OCI GoldenGate is available.

See [GoldenGate API Endpoints](#) to find the API endpoints for each region.

Service Limits

OCI GoldenGate limits you to 20 deployments and 100 registered databases per region.

Compartment Quotas

Creating a quota limit lets you limit the number of deployment resources in a compartment.

For example:

```
set goldengate quota deployment-count to 5 in compartment  
<compartment_name>
```

What's next?

Ensure that you:

- Create Oracle Cloud resources for OCI GoldenGate
- Configure identity domains, if using an OCI Identity Access Management (OCI IAM) enabled tenancy.

What's supported

Learn about OCI GoldenGate deployment types and the connection types they support.

Deployment types

When you create a deployment, you select a deployment type for your specific data management needs:

- Data replication
- Data transforms
- Stream analytics

The following lists detail the source and target technologies supported by each deployment type.

Supported connection types for data replication

A connection contains the network connectivity details to a data source or target.

Table 3-1 Connections supported by each deployment technology type



Technology type	Connection types supported
Oracle Database	<ul style="list-style-type: none"> • Oracle Database 11.2.0.4, 12.1.0.2, and higher • Oracle Exadata 11.2.0.4, 12.1.0.2, and higher • Oracle Exadata Cloud Service • Oracle Exadata Cloud@Customer (remote capture and delivery) • Oracle Autonomous Transaction Processing • Oracle Autonomous Data Warehouse • Amazon RDS for Oracle 19 and higher
	<div style="border: 1px solid #0070C0; padding: 10px; background-color: #E6F2FF;"> <p> Note:</p> <p>Ensure that you apply the latest patches for the databases you use with OCI GoldenGate as recommended by Oracle Support.</p> <ul style="list-style-type: none"> • Recommended patches for Oracle 11g • Recommended patches for Oracle 12c or higher </div>
Big Data	<p>The Big Data deployment type is available in OCI GoldenGate build 21.6.0.0.2_220714.0903_616 and higher.</p> <ul style="list-style-type: none"> • Oracle Autonomous Database (target only) • OCI Object Storage (target only) • OCI Streaming • Oracle Oracle Autonomous JSON Database (target only) • Oracle NoSQL (target only) • Oracle Oracle Weblogic JMS • Apache Kafka 2.4, 2.5, 2.6, 2.7, 2.8, 3.0, 3.1, 3.2 • Confluent Kafka, with or without Confluent Schema Registry, 7.1 or earlier • Snowflake (target only) • MongoDB • Amazon MSK • Amazon S3 (target only) • Amazon Redshift (target only) • Amazon Kinesis (target only) • Azure Data Lake Storage / Azure BLOB Storage (target only) • Azure Synapse Analytics (target only) • Azure Cosmos DB for MongoDB (target only) • Azure Event Hubs • Google BigQuery (target only) • Google Cloud Storage (target only) • Elasticsearch Server (target only) • Hadoop Distributed File System (HDFS) (target only) • Redis

Table 3-1 (Cont.) Connections supported by each deployment technology type

Technology type	Connection types supported
Microsoft SQL Server	<p>The Microsoft SQL Server deployment type is available in OCI GoldenGate build 21.9.0.0.0_230120.0600_716 and higher.</p> <ul style="list-style-type: none"> • Azure SQL Database (target only) • Azure SQL Managed Instance • Microsoft SQL Server 2012, 2014, 2016, 2017, 2019, 2022 <div style="border: 1px solid #0070C0; padding: 10px; margin: 10px 0;"> <p> Note:</p> <p>Microsoft SQL Server 2022 is only supported in builds _802 and higher</p> </div> <ul style="list-style-type: none"> • Amazon RDS for SQL Server 2012, 2014, 2016, 2017, 2019 • Google Cloud SQL for SQL Server
MySQL	<p>The MySQL deployment type is available in OCI GoldenGate build 21.7.0.0.0_220820.1908_616 and higher.</p> <ul style="list-style-type: none"> • MySQL Database Server 5.7, 8.0 • OCI MySQL Heatwave 8.0 • Amazon RDS for MySQL 5.7, 8.0 • Azure Database for MySQL 5.7, 8.0 • Amazon Aurora MySQL 5.7, 8.0 • Google Cloud SQL for MySQL 8.0 • Amazon RDS for MariaDB 10.4, 10.5 • MariaDB 10.4, 10.5, 10.6, 10.7, 10.8, 10.9 • SingleStoreDB • SingleStoreDB Cloud
PostgreSQL	<p>The PostgreSQL deployment type is available in OCI GoldenGate build 21.7.0.0.0_220731.2140_663 and higher.</p> <ul style="list-style-type: none"> • PostgreSQL Server 10, 11, 12, 13, 14, 15 • Amazon Aurora PostgreSQL 10, 11, 12, 13 • Amazon RDS PostgreSQL 10, 11, 12, 13, 14 • Azure Database for PostgreSQL 10, 11, 12, 13 • Google Cloud SQL for PostgreSQL 10, 11, 12, 13, 14
Generic	<ul style="list-style-type: none"> • GoldenGate • Generic connection

For more information on build versions, see OCI GoldenGate Versions.

Learn more

For more information about what's supported and not supported, see:

- Understanding what's supported for Oracle Database
- Understanding what's supported for MySQL
- Understanding what's supported for PostgreSQL
- Understanding what's supported for SQL Server

Supported connection types for Data Transforms

See Supported Connection Types for a full list of connections that you can use with OCI GoldenGate Data Transforms.

 **Tip:**

For each connection type used with OCI GoldenGate Data Transforms, you must create a Generic connection, and assign the connection to the Data Transforms deployment.

After assigning the Generic connection(s) to the Data Transforms deployment, launch the console from the deployment details page and log in. In the Data Transforms deployment console, create a connection for each data source using the Host name(s) provided for the Generic connection(s).

Supported connection types for Stream Analytics

Table 3-2 Connections supported by Stream Analytics

Source or Target	Connection types supported
Sources	<ul style="list-style-type: none"> • Oracle Autonomous Transaction Processing • Oracle Autonomous Data Warehouse • Oracle Database • OCI Streaming • MySQL Database Server 5.7, 8.0 • Apache Kafka • Confluent Kafka • Coherence, Ignite, and Java Message Server within the Stream Analytics console • Oracle GoldenGate
Targets	<ul style="list-style-type: none"> • Oracle Autonomous Transaction Processing • Oracle Autonomous Data Warehouse • Oracle Database • OCI Object Storage • OCI Streaming • Apache Kafka • Amazon S3, Azure Data Lake Storage, Coherence, Hadoop File Storage (HDFS), Ignite, Java Message Server, and MongoDB within the Stream Analytics console

For more information on build versions, see OCI GoldenGate Versions.

OCPU management and billing

Learn about Oracle Compute Units (OCPU) and how you're billed for usage.

**Note:**

Customers using features in released in Limited Availability, you're not billed for usage of these resources, however, you must be aware of how you'll be billed when Limited Availability ends.

Metering and billing

Metering and billing for OCI GoldenGate is based on the number of OCPUs the service uses per minute.

When you create an OCI GoldenGate deployment, you select the number of OCPUs between 1 and 24 that your deployment will use. Each OCPU allocates an additional 16 GB memory.

You can enable auto scaling, which allows the service to scale up to three times the OCPU Count you specified upon creation. When you enable auto scaling, you are billed for the actual average number of OCPUs consumed per hour.

For example, if you specify 3 as the base OCPU and enable Auto Scale, then the total OCPUs that can be used is 9. When the OCPU Utilization is greater than 33.333% of 9 OCPUs, you are billed for the integer value over 33.333%, which is 4 OCPUs.

OCPU Utilization Greater than	OCPU Utilization Less than or equal to	Billed for
0	33.333%	3 OCPUs
33.333%	44.444%	4 OCPUs
44.444%	55.555%	5 OCPUs
55.555%	66.666%	6 OCPUs
66.666%	77.777%	7 OCPUs
77.777%	88.888%	8 OCPUs
88.888%	100%	9 OCPUs

OCI GoldenGate cannot scale over 24 OCPUs, which is the maximum number of OCPUs. For example, if you select 9 OCPUs as your base OCPU Count and enable Auto Scale, the service will scale up to 24 OCPUs. You can enable or disable auto scaling at any time.

You can monitor the OCPU consumption and memory of a deployment in the Metrics section of a deployment's Detail page.

About OCPU utilization

Before you create an OCI GoldenGate deployment, it's important to understand how the OCPU shape affects the OCI GoldenGate user experience.

When you create an OCI GoldenGate deployment, you must select how many OCPUs the deployment will use. Here are some things to consider when making your selection:

- 1 OCPU shapes are highly recommended for non-production environments.

- Enabling Auto Scale ensures additional OCPUs are available when they're needed. You're only billed when the additional OCPU cycles are consumed.
- Extract and Replicat are high priority processes and take precedence in OCPU cycles.
- OCI GoldenGate Deployment Console also consumes OCPU cycles. If the current OCPU utilization is near 100%, then the OCI GoldenGate Deployment Console may be very slow to return or not return at all.
- The OCI GoldenGate Backup process also consume OCPU cycles. If the current OCPU utilization is near 100%, the backup process can take a long time to complete.
- If an OCI GoldenGate deployment is running at or near 100% OCPU utilization, consider changing the base OCPU count or enabling Auto Scale, if it's not already enabled.
- You can monitor OCPU Utilization on the OCI GoldenGate Deployment Details page and the Metrics Explorer in the OCI Console.
- You can [set alarms in the OCI Console](#) to notify you when OCPU utilization reaches certain levels.

Metering and billing for Stream Analytics deployments

Ensure that you review the information in Metering and billing for OCI GoldenGate deployments about Oracle Compute Unit (OCPU) selection and scaling.

OCI GoldenGate Stream Analytics OCPU usage is calculated based on the following factors:

- Stream Analytics console
- Number of Streaming pipelines
- Ignite cluster
- GoldenGate Big Data cluster

Before calculating the number of OCPUs you need, let's first review how many compute units each Stream Analytics resource requires. 1 OCPU is equal to 2 compute units (vCPUs). 1 vCPU is equal to 1000 millicores (1000m).

The following table lists example Stream Analytics pipeline settings and the calculated number of OCPUs required.

Pipeline	Driver	Executor	Total vCPUs	OCPUs billed
Pipeline A	500m	1 x 500m	1000m	1
Pipeline B	500m	2 x 500m	1500m	1
Pipeline C	500m	4 x 500m	2500m	2
Pipeline D	600m	2 x 700m	2000m	1
Pipeline E	1000m	2 x 1000m	3000m	2

You can configure the Driver and Executor settings as needed for each pipeline in the Stream Analytics console.

The following table lists example Stream Analytics resource configurations based on the number of pipelines (from the above table) and the calculated number of OCPUs required.

Stream Analytics console	Number of pipelines	Streaming pipelines	Ignite cluster	GoldenGate for Big Data cluster	OCPUs billed
1000m	1 x Pipeline A	1000m	0	0	1
1000m	3 x Pipeline A	3000m	0	0	2
1000m	1 x Pipeline B	1500m	0	0	2
1000m	1 x Pipeline B	1500m	2 x 500m	500m	2
1000m	1 x Pipeline A 1 x Pipeline B	2500m	2 x 500m	500m	3
1000m	2 x Pipeline A 1 x Pipeline B	3500m	2 x 500m	500m	3

The Stream Analytics console requires 1000m. Each streaming pipeline requires additional millicores depending on their settings. The Ignite cluster, if activated, requires a minimum of 2 cluster instances. You can configure the millicore limit for both Ignite and GoldenGate Big Data clusters in the Stream Analytics console. When added together, you can determine the total number of OCPUs that you need to select when creating your Stream Analytics deployment.

If you're unsure, you can start with 2 or 3 OCPUs, and then review the OCPU consumption metrics on the deployment details page and adjust accordingly.

Integrated services

Learn about other Oracle Cloud Infrastructure services that are integrated with OCI GoldenGate.

IAM

OCI GoldenGate integrates with the Identity and Access Management (IAM) service for authentication and authorization for the Console, SDK, CLI, and REST API, but not the OCI GoldenGate Deployment Console. To access the OCI GoldenGate Deployment Console, use the username and password that you specify when you create the deployment.

To learn more about the OCI GoldenGate Deployment Console, see [Explore the deployment console](#).

To learn more about IAM, see [IAM Overview](#).

Events

Oracle Cloud Infrastructure Events lets you create automation based on the state changes of OCI GoldenGate resources.

The following OCI GoldenGate resources emit events:

- Deployments
- Connections
- Deployment backups

- Upgrade notifications

Learn more about GoldenGate Events.

Work Requests

OCI GoldenGate is not integrated with the common Work Requests API. The Oracle Cloud Infrastructure GoldenGate service uses its own API for Work Requests. See [Work Request Reference](#).

Monitoring

Oracle Cloud Infrastructure Monitoring lets you actively and passively monitor your Oracle Cloud Infrastructure GoldenGate resources and alarms. GoldenGate Metrics capture CPU utilization, OCPU consumption, memory utilization, deployment health, and inbound and outbound lag. You can view these metrics using the Monitoring service, or on your Deployment Details page.

Example OCI GoldenGate topologies

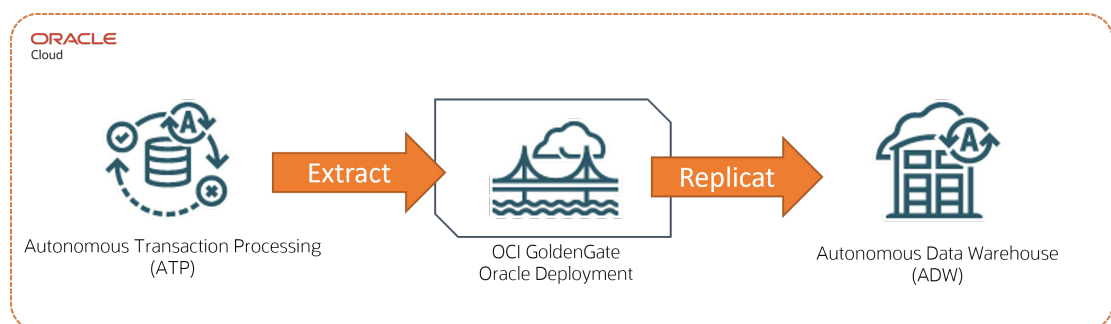
Before you create your OCI GoldenGate deployments, review these sample topologies to help you plan for the number of resources your solution needs.

How many resources do I need?

Deployments

To determine how many deployments your solution needs, consider the types of technologies you're replicating data between.

For example, if your source and target databases are Oracle Autonomous Databases, then you need only one Oracle deployment type.



See Replicate data between cloud databases in the same region.

If you're replicating data between two different technologies, then you need two OCI GoldenGate deployments. For example, if your source database is a MySQL database type, and your target is a Big Data type, then you must:

- Create a MySQL deployment for your MySQL source
- Create a Big Data deployment for your Big Data target

This solution also requires a Distribution Path. See the following examples for details.

Connections

You must create a connection for each source and target technology, and then assign the connections to the appropriate deployment. Using the MySQL to Big Data example above, you must:


- Create a connection to your source MySQL database, and then assign it to the MySQL deployment
- Create a connection to your target Big Data technology, and then assign it to the Big Data deployment

Note:

If your target deployment doesn't have a public endpoint, then you must also create a GoldenGate connection, and then assign this connection to the source deployment.

Example: Azure SQL Managed Instance to Autonomous Transaction Processing

In this example, Azure SQL Managed Instance is the source technology and Autonomous Transaction Processing (ATP) is the target.




For this replication scenario, you need:

- Two deployments:
 - A Microsoft SQL Server deployment for the source database
 - An Oracle deployment for the target database
- Connections:
 - A connection to Azure SQL Managed Instance, and then assigned to the Microsoft SQL Server deployment
 - A connection to Autonomous Transaction Processing, and then assigned to the Oracle deployment
 - If your target deployment doesn't have a public endpoint, create an Oracle GoldenGate connection and then assign it to the source deployment.
- Processes:
 - An Extract process created in the source deployment
 - A Distribution Path created in the source deployment
 - A Replicat created in the target deployment

This replication scenario is available as quickstart.

Example: Autonomous Transaction Processing to Apache Kafka

In this example, Autonomous Transaction Processing (ATP) is the source technology and Apache Kafka is the target.




For this replication scenario, you need:

- Two deployments:
 - An Oracle deployment for the source database
 - A Big Data deployment for the target technology
- Connections:
 - A connection to Autonomous Transaction Processing, and then assigned to the Oracle deployment
 - A connection to Apache Kafka, and then assigned to the Big Data deployment
 - If your target deployment doesn't have a public endpoint, create an Oracle GoldenGate connection and then assign it to the source deployment.
- Processes:
 - An Extract process created in the source deployment
 - A Distribution Path created in the source deployment
 - A Replicat created in the target deployment

This replication scenario is available as a quickstart.

Example: PostgreSQL to Autonomous Transaction Processing

In this example, PostgreSQL is the source technology and Autonomous Transaction Processing (ATP) is the target.



For this replication scenario, you need:


- Two deployments:
 - A PostgreSQL deployment for the source database
 - An Oracle deployment for the target technology

- Connections:
 - A connection to PostgreSQL, and then assigned to the PostgreSQL deployment
 - A connection to Autonomous Transaction Processing, and then assigned to the Oracle deployment
 - If your target deployment doesn't have a public endpoint, create an Oracle GoldenGate connection and then assign it to the source deployment.
- Processes:
 - An Extract process created in the source deployment
 - A Distribution Path created in the source deployment
 - A Replicat created in the target deployment

This replication scenario is available as a quickstart.

Example: PostgreSQL to MySQL

In this example, PostgreSQL is the source technology and MySQL is the target.




For this replication scenario, you need:

- Two deployments:
 - A PostgreSQL deployment for the source database
 - A MySQL deployment for the target technology
- Connections:
 - A connection to PostgreSQL, and then assigned to the PostgreSQL deployment
 - A connection to MySQL, and then assigned to the MySQL deployment
 - If your target deployment doesn't have a public endpoint, create an Oracle GoldenGate connection and then assign it to the source deployment.
- Processes:
 - An Extract process created in the source deployment
 - A Distribution Path created in the source deployment
 - A Replicat created in the target deployment

This replication scenario is available as a quickstart.

Example: PostgreSQL to Snowflake

In this example, PostgreSQL is the source technology and Snowflake is the target.



For this replication scenario, you need:

- Two deployments:
 - A PostgreSQL deployment for the source database
 - A Big Data deployment for the target technology
- Connections:
 - A connection to PostgreSQL, and then assigned to the PostgreSQL deployment
 - A connection to Snowflake, and then assigned to the Big Data deployment
 - If your target deployment doesn't have a public endpoint, create an Oracle GoldenGate connection and then assign it to the source deployment.
- Processes:
 - An Extract process created in the source deployment
 - A Distribution Path created in the source deployment
 - A Replicat created in the target deployment

This replication scenario is available as a quickstart.

4

Connect

Learn to create the connections to source and target technologies.

Articles in this section:

- [Explore connections](#)
- [Oracle Database connections](#)
- [Big Data connections](#)
- [MySQL connections](#)
- [PostgreSQL connections](#)
- [SQL Server connections](#)
- [Create a Generic connection](#)
- [Connect to GoldenGate Distribution and Receiver Paths](#)

Explore connections

Learn about the different types of connections you can use with OCI GoldenGate and how to create them.

What is a connection?

A connection contains the connectivity details for a data source or target. Connections support database and non-database technologies such as Oracle, MySQL, PostgreSQL, Microsoft SQL Server, Kafka, Oracle Object Storage, OCI Streaming, and Oracle GoldenGate Distribution and Receiver Paths, to name a few.

WARNING:

You must only create and edit connections in the Oracle Cloud console. Refrain from creating or editing connections in the Credentials screen of the deployment console. Updates are automatically synced to the deployment from the Oracle Cloud console.

Before you create connections, it's recommended that you review OCI GoldenGate connectivity to learn how OCI GoldenGate connects to your sources and targets.

After you create a connection, you must then assign it to the deployment you want to use it with. Keep in mind that certain deployment types support specific connection types. You can assign a connection to a deployment that resides in a different compartment. To view the assignment from the connection details page, you must first change the compartment view.

Learn more about managing connections.

Supported connection types for data replication

A connection contains the network connectivity details to a data source or target.

Table 4-1 Connections supported by each deployment technology type



Technology type	Connection types supported
Oracle Database	<ul style="list-style-type: none"> • Oracle Database 11.2.0.4, 12.1.0.2, and higher • Oracle Exadata 11.2.0.4, 12.1.0.2, and higher • Oracle Exadata Cloud Service • Oracle Exadata Cloud@Customer (remote capture and delivery) • Oracle Autonomous Transaction Processing • Oracle Autonomous Data Warehouse • Amazon RDS for Oracle 19 and higher
	<div style="border: 1px solid #0070C0; padding: 10px; background-color: #E6F2FF;"> <p> Note:</p> <p>Ensure that you apply the latest patches for the databases you use with OCI GoldenGate as recommended by Oracle Support.</p> <ul style="list-style-type: none"> • Recommended patches for Oracle 11g • Recommended patches for Oracle 12c or higher </div>
Big Data	<p>The Big Data deployment type is available in OCI GoldenGate build 21.6.0.0.2_220714.0903_616 and higher.</p> <ul style="list-style-type: none"> • Oracle Autonomous Database (target only) • OCI Object Storage (target only) • OCI Streaming • Oracle Oracle Autonomous JSON Database (target only) • Oracle NoSQL (target only) • Oracle Oracle Weblogic JMS • Apache Kafka 2.4, 2.5, 2.6, 2.7, 2.8, 3.0, 3.1, 3.2 • Confluent Kafka, with or without Confluent Schema Registry, 7.1 or earlier • Snowflake (target only) • MongoDB • Amazon MSK • Amazon S3 (target only) • Amazon Redshift (target only) • Amazon Kinesis (target only) • Azure Data Lake Storage / Azure BLOB Storage (target only) • Azure Synapse Analytics (target only) • Azure Cosmos DB for MongoDB (target only) • Azure Event Hubs • Google BigQuery (target only) • Google Cloud Storage (target only) • Elasticsearch Server (target only) • Hadoop Distributed File System (HDFS) (target only) • Redis

Table 4-1 (Cont.) Connections supported by each deployment technology type

Technology type	Connection types supported
Microsoft SQL Server	<p>The Microsoft SQL Server deployment type is available in OCI GoldenGate build 21.9.0.0.0_230120.0600_716 and higher.</p> <ul style="list-style-type: none"> • Azure SQL Database (target only) • Azure SQL Managed Instance • Microsoft SQL Server 2012, 2014, 2016, 2017, 2019, 2022 <div style="border: 1px solid #0070C0; padding: 10px; margin: 10px 0;"> <p> Note:</p> <p>Microsoft SQL Server 2022 is only supported in builds _802 and higher</p> </div> <ul style="list-style-type: none"> • Amazon RDS for SQL Server 2012, 2014, 2016, 2017, 2019 • Google Cloud SQL for SQL Server
MySQL	<p>The MySQL deployment type is available in OCI GoldenGate build 21.7.0.0.0_220820.1908_616 and higher.</p> <ul style="list-style-type: none"> • MySQL Database Server 5.7, 8.0 • OCI MySQL Heatwave 8.0 • Amazon RDS for MySQL 5.7, 8.0 • Azure Database for MySQL 5.7, 8.0 • Amazon Aurora MySQL 5.7, 8.0 • Google Cloud SQL for MySQL 8.0 • Amazon RDS for MariaDB 10.4, 10.5 • MariaDB 10.4, 10.5, 10.6, 10.7, 10.8, 10.9 • SingleStoreDB • SingleStoreDB Cloud
PostgreSQL	<p>The PostgreSQL deployment type is available in OCI GoldenGate build 21.7.0.0.0_220731.2140_663 and higher.</p> <ul style="list-style-type: none"> • PostgreSQL Server 10, 11, 12, 13, 14, 15 • Amazon Aurora PostgreSQL 10, 11, 12, 13 • Amazon RDS PostgreSQL 10, 11, 12, 13, 14 • Azure Database for PostgreSQL 10, 11, 12, 13 • Google Cloud SQL for PostgreSQL 10, 11, 12, 13, 14
Generic	<ul style="list-style-type: none"> • GoldenGate • Generic connection

For more information on build versions, see OCI GoldenGate Versions.

Supported connection types for Data Transforms

See Supported Connection Types for a full list of connections that you can use with OCI GoldenGate Data Transforms.



Tip:

For each connection type used with OCI GoldenGate Data Transforms, you must create a Generic connection, and assign the connection to the Data Transforms deployment.

After assigning the Generic connection(s) to the Data Transforms deployment, launch the console from the deployment details page and log in. In the Data Transforms deployment console, create a connection for each data source using the Host name(s) provided for the Generic connection(s).

Supported connection types for Stream Analytics

Table 4-2 Connections supported by Stream Analytics

Source or Target	Connection types supported
Sources	<ul style="list-style-type: none"> • Oracle Autonomous Transaction Processing • Oracle Autonomous Data Warehouse • Oracle Database • OCI Streaming • MySQL Database Server 5.7, 8.0 • Apache Kafka • Confluent Kafka • Coherence, Ignite, and Java Message Server within the Stream Analytics console • Oracle GoldenGate
Targets	<ul style="list-style-type: none"> • Oracle Autonomous Transaction Processing • Oracle Autonomous Data Warehouse • Oracle Database • OCI Object Storage • OCI Streaming • Apache Kafka • Amazon S3, Azure Data Lake Storage, Coherence, Hadoop File Storage (HDFS), Ignite, Java Message Server, and MongoDB within the Stream Analytics console

For more information on build versions, see OCI GoldenGate Versions.

Oracle Database connections

Learn to create connections to Oracle Databases including:

- [Connect to Oracle Autonomous Databases](#)
- [Connect to Oracle Database](#)
- [Connect to Oracle Exadata](#)
- [Connect to Amazon RDS Oracle Database](#)

Connect to Oracle Autonomous Database

Learn to create a connection to Oracle Autonomous Transaction Processing or Autonomous Data Warehouse to use as sources and targets for OCI GoldenGate.



Note:

If creating a connection to an Autonomous Database Dedicated workload type, skip to the known issue entry in this article.

Connect to Autonomous Database Shared

To create an Autonomous Database Shared connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Oracle Autonomous Database**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. On the **Connection Details** fields, for Database details:
 - Choose **Select database** to select from a list of existing Autonomous Databases in the selected compartment, and then enter the `ggadmin` password. Click **Change compartment** to choose a database in a different compartment.



Note:

When you select an existing Autonomous Database, a private endpoint is created automatically.

- Choose **Enter database information** and then manually complete the following fields:
 - a. For **Database username**, enter the username to connect to the database with.
 - b. For **Database password**, enter the password associated with the username entered in the previous step.
 - c. For **Upload wallet**, drag-and-drop or upload the database wallet.zip for **Database wallet**.
 - d. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
Then, select a **Session mode**:
 - * **Direct**, to use the local listener running on a single database node, and then select your subnet.
 - * **Redirect**, to use the SCAN listener used in Oracle Real Application Cluster (RAC) deployments, and then select your subnet.

 **Tip:**

Subnet is only required when you enter the database details manually, and the database is public. Private databases already have their own subnet.

 **Note:**

Learn more about Oracle GoldenGate connectivity.

6. Click **Create**.

After the connection is created, it appears in the Connections list. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Known issues

Create connection for Autonomous Database Dedicated must use Oracle Database type

If creating a connection for an *Autonomous Database Dedicated* instance, then you must select Oracle Database as the type, instead of Oracle Autonomous Database.

To create a connection for Autonomous Database Dedicated:

1. On the Connections page, click **Create connection**.

2. In the Create connections panel, enter a name for the connection, and optionally, a description.
3. Select a compartment in which to create the connection.
4. From the Type dropdown, select **Oracle Database**.
5. Click **Next**.
6. On the Connection details screen, for **Database details**, select **Enter database information**.
7. For **Database connection string**, enter the TCP connection string. You can find the connection string in the tnsnames.ora file downloaded from the wallet package on the Autonomous Database Dedicated details page. Ensure you use the "_low" connection string.
8. Enter the database username and password.
9. **Do not** upload the database wallet.
10. In the Network connectivity section:
 - a. Select **Dedicated endpoint**.
 - b. For **Session mode**, select **Redirect**.
11. Click **Create**.

Action Required for Autonomous Databases that Use mTLS Authentication

When an Autonomous Database wallet is rotated, the OCI GoldenGate connection to this database must be refreshed to retrieve the latest wallet information.

For more information see, [My Oracle Support \(MOS\) Document 2911553.1](#).

To refresh an Autonomous Database connection: Edit and save the connection to the Autonomous Database (Autonomous Transaction Processing or Autonomous Datawarehouse). Saving the connection automatically downloads and refreshes the wallet. No other changes to the connection is needed.

To verify:

1. Launch the deployment console for a deployment that uses the Autonomous Database connection.
2. In the deployment console, open the navigation menu, and then click **Configuration**.
3. On the Credentials screen, observe the Autonomous Database connection string. Before the wallet is refreshed, the connection string looks like the following:

```
ggadmin@(DESCRIPTION=(TRANSPORT_CONNECT_TIMEOUT=3) (CONNECT_TIMEOUT=60)
(RECV_TIMEOUT=120) (retry_count=20) (retry_delay=3) (address=(protocol=tcps)
(port=1522) (host=adb.us-phoenix-1.oraclecloud.com) )
(CONNECT_DATA=(COLOCATION_TAG=ogginstance) (FAILOVER_MODE=(TYPE=SESSION)
(METHOD=BASIC) (OVERRIDE=TRUE)) (service_name=<adb-
servicename>_low.adb.oraclecloud.com))
(security=(MY_WALLET_DIRECTORY="/u02/connections/
ocidl.goldengateconnection.oc1.phx.<ocid>/wallet")
(SSL_SERVER_DN_MATCH=TRUE) (ssl_server_cert_dn="CN=adwc.uscom-
east-1.oraclecloud.com,
OU=Oracle BMCS US, O=Oracle Corporation, L=Redwood City,
```



```
ST=California,  
  C=US"))))
```

After the wallet is refreshed, the connection string is updated to look like the following:

```
ggadmin@ (DESCRIPTION= (TRANSPORT_CONNECT_TIMEOUT=3)  
 (CONNECT_TIMEOUT=60) (RECV_TIMEOUT=120) (retry_count=20)  
 (retry_delay=3) (address=(protocol=tcps) (port=1522) (host=adb.us-  
phoenix-1.oraclecloud.com))  
 (CONNECT_DATA=(COLOCATION_TAG=ogginstance)  
 (FAILOVER_MODE=(TYPE=SESSION) (METHOD=BASIC) (OVERRIDE=TRUE))  
 (service_name=<adb-servicename>_low.adb.oraclecloud.com))  
 (security=(MY_WALLET_DIRECTORY="/u02/connections/  
ocid1.goldengateconnection.oc1.phx.<ocid>/wallet")  
 (SSL_SERVER_DN_MATCH=TRUE) (ssl_server_dn_match=yes)))
```

Creating a connection to Oracle Database

Learn how to create a connection to an Oracle Database to use as a source or target.

Create the connection

To create an Oracle Database connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Oracle Database**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection Details** fields as follows:

- a. Choose **Select database** to select an existing database in the selected compartment, and then complete the rest of the fields as needed. Click **Change Compartment** to select a database in a different compartment. Complete the following fields:
- b. Choose **Enter database information** to enter the database details manually:
 - i. (Optional) For **Database connection string**, enter the database's connection string.

 **Note:**

To find the connection string, click **DB Connection** on the database's Details page.

- ii. For **Database username**, enter the username to connect to the database with.
- iii. For **Database password**, enter the password associated to the username provided in the previous step.

 **Note:**

If using Oracle Database, ensure that you use the CDB user to capture data from PDBs. See *Configuring Oracle GoldenGate in a Multitenant Container Database*.

- iv. (Optional) For **Database wallet**, drag and drop or select the wallet.zip for this database.

 **Note:**

The wallet.zip must contain the `cwallet.sso` and `tnsnames.ora` files. Within the `tnsnames.ora`, one of the entries must have the suffix, `_low`, and description, security, and `connect_data` parameter value pairs in all lowercase letters. For example:

```
wdb19c_low =
  (description=(address=(protocol=tcps) (host=<hostname>)
(port=<port_number>)) (connect_data=(server=dedicated)
(service_name=<service_name>))
(security=(ssl_server_dn_match=true)
(ssl_server_cert_dn="cn=<cert_name>")))
```

- c. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.

 **Note:**

Select this option for on premise Oracle Databases not available over the public internet.

Then, select a **Session mode**:

- **Direct**, to use the local listener running on a single database node, and then select your subnet.
- **Redirect**, to use the SCAN listener used in Oracle Real Application Cluster (RAC) deployments, and then select your subnet.

 **Tip:**

Subnet is only required when you enter the database details manually, and the database is public. Private databases already have their own subnet.

 **Note:**

Learn more about Oracle GoldenGate connectivity.

6. Click **Create**.

After the connection is created, it appears in the Connections list. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Known issues

Network timeout affects database connections using private endpoints.

If you're using a private endpoint to connect to a database, then you may encounter network timeouts when starting or stopping Extract processes.

Workaround: You can do one of the following:

- Apply the latest patches from your deployment details page. In the **Deployment Information** section, under **GoldenGate**, for **Version**, click **Upgrade**.
- If you're unable to apply the latest patches at this time, you can update the connection string to include `EXPIRE_TIME=1`. By default, you may have an EZ connection string in Oracle GoldenGate. This connection string needs to be updated in the Oracle GoldenGate Credential to a long connection string as follows:

```
<username>@//<hostname>:1521/<service_name>
<username> @(DESCRIPTION = (EXPIRE_TIME=1)(ADDRESS_LIST = (ADDRESS
= (COMMUNITY = tcp)(PROTOCOL = TCP)(Host = <hostname>)(Port =
1521))) (CONNECT_DATA = (SERVICE_NAME = <service_name>)))
```

SCAN Proxy doesn't support TLS

While OCI GoldenGate supports Oracle Single Client Access Name (SCAN) hosts and IPs, the SCAN proxy does not support TLS.

Workaround: You can connect to a RAC database using the Database Node IP.

Connect to Oracle Exadata

Learn to create a connection to an Oracle Exadata database to use as an OCI GoldenGate source or target.

To create an Oracle Exadata connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Oracle Exadata**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. Choose **Select database** to select an existing database in the selected compartment, and then complete the rest of the fields as needed. Click **Change compartment** to select a database in a different compartment.
 - b. Choose **Enter database information** to enter the database details manually:
 - i. (Optional) For **Database connection string**, enter the database's connection string.

 **Note:**

To find the connection string, click **Database connection** on the database's Details page.

- ii. For **Database username**, enter the username to connect to the database with.
 - iii. For **Database password**, enter the password associated to the username provided in the previous step.
 - iv. (Optional) For **Database wallet**, drag and drop or select the wallet.zip for this database.
- c. For **Network connectivity**, select a **Traffic routing method**:
- **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.

 **Note:**

Select this option for on premise Oracle Databases not available over the public internet.

Then, select a **Session mode**:

- **Direct**, to use the local listener running on a single database node, and then select your subnet.
- **Redirect**, to use the SCAN listener used in Oracle Real Application Cluster (RAC) deployments, and then select your subnet.

 **Tip:**

Subnet is only required when you enter the database details manually, and the database is public. Private databases already have their own subnet.

 **Note:**

Learn more about Oracle GoldenGate connectivity.

6. Click **Create**.

After the connection is created, it appears in the Connections list. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Connect to Amazon RDS Oracle Database

Learn to create a connection to Amazon RDS for Oracle database to use as a source or target for OCI GoldenGate.

To create an Amazon RDS for Oracle database connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.

2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Amazon RDS for Oracle**.
 - e. (Optional) Click **Show Advanced Options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection Details** fields as follows:
 - a. (Optional) For **Database connection string**, enter the database's connection string.
 - b. For **Database username**, enter the username to connect to the database with.
 - c. For **Database password**, enter the password associated to the username provided in the previous step.
 - d. (Optional) For **Database wallet**, drag and drop or select the wallet.zip for this database.
 - e. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.

 **Note:**

Select this option if your Azure SQL Managed Instance doesn't have a public endpoint.

If selected:

- Select the subnet
- (Optional) Enter the Private IP only if the hostname is not resolvable from your subnet or if it uses SSL/TLS.

 **Note:**

OCI GoldenGate rewrites the private IP in the format, `ip-10-0-0-0.ociggsvc.oracle.vcn.com`.

 **Note:**

Learn more about Oracle GoldenGate connectivity.

6. Click **Create**.

After the connection is created, it appears in the Connections list where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Big Data connections

Learn to create connections to Big Data technologies, including:

- [Connect to Oracle Autonomous JSON Database](#)
- [Connect to OCI Object Storage](#)
- [Connect to OCI Streaming](#)
- [Connect to Oracle NoSQL](#)
- [Connect to Oracle Weblogic JMS](#)
- [Connect to Amazon Kinesis](#)
- [Connect to Amazon Redshift](#)
- [Connect to Amazon S3](#)
- [Connect to Apache Kafka](#)
- [Connect to Confluent Kafka](#)
- [Connect to Confluent Schema Registry](#)
- [Connect to Azure Cosmos DB for MongoDB](#)
- [Connect to Azure Data Lake Storage](#)
- [Connect to Azure Synapse Analytics](#)
- [Connect to Azure Event Hubs](#)
- [Connect to MongoDB](#)
- [Connect to Redis](#)
- [Connect to Snowflake](#)
- [Connect to Google Cloud Storage](#)
- [Connect to Google BigQuery](#)
- [Connect to Elasticsearch Server](#)
- [Connect to Hadoop Distributed File System](#)

Create a connection to Oracle Autonomous JSON Database

Learn to create a connection to an Oracle Autonomous JSON Database to use as a target in a replication process.

To create an Oracle Autonomous JSON Database connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. Select **Oracle Autonomous JSON Database** from the **Type** dropdown.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. In the Database details section, choose:
 - **Select database** to select from a list of existing Autonomous JSON Databases in the compartment. Click **Change Compartment** to select a database in a different compartment.
 - **Enter database information** and then manually enter the **Connection String**.
 - b. Enter the Username and Password for the Oracle Autonomous JSON Database.
 - c. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

After creating the connection, it appears in the Connections list where you can view its details. Ensure that you assign the connection to a Big Data deployment to use it as a target in a replication.

Connect to OCI Object Storage

Learn to create a connection to OCI Object Storage to use as a target with OCI GoldenGate

Create the connection

To create an OCI Object Storage connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Technology Type**, select **OCI Object Storage**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Tenancy**, select one of the following options:
 - Use current tenancy
 - Specify another tenancy, and then enter the **Tenancy OCID**
 - b. For **Region**, select the OCI Object Storage region.
 - c. For **User**, select one of the following options:
 - Use current user
 - Specify another user, and then enter the **User OCID**.
 - d. For **Private key file**, drag and drop or select the Privacy enhanced mail (PEM) file.

 **Tip:**

You can create and download the API key on your user details page, and then upload it here.

- e. For **Public key fingerprint**, enter the corresponding public key fingerprint.
- f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

After the connection is created, it appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a target in a replication.

Known Issues

User OCID Mismatch in OCI Object Storage connection (Federated users only)

If a federated user selects Use current user when creating an OCI Object Storage connection, their OCID doesn't match the OCID picked up by the system.

Workaround: When you create an OCI Object Storage connection, ensure that you choose **Specify another user**, and then enter the federated user's OCID.

To find the user OCID, click **Profile** in the Oracle Cloud console global header, and then select the user name. On the User Details page, under User Information, click **Show** for OCID.

Connect to OCI Streaming

Learn to create a connection to OCI Streaming to use as a source or target with OCI GoldenGate

Before you begin

1. Obtain the Stream Pool username:
 - a. From the Oracle Cloud console navigation menu, select **Streaming**, and then **Stream Pools**.
 - b. On the Stream Pools page, select your pool to view its details.
 - c. On the Stream Pool details page, under **Resources**, click **Kafka Connection Settings**.
 - d. Copy the username for SASL Connection Strings.
2. Create an Auth token:
 - a. In the Oracle Cloud console global header, click **Profile**, and then select **User settings**.

- b. On the User Details page, under **Resources**, click **Auth Tokens**, and then click **Generate Token**.
- c. In the Generate Token dialog, enter a description, and then click **Generate Token**.
- d. Copy the auth token from the dialog to a secure location from where you can retrieve it later, and then click **Close**.

The Stream Pool username and Auth token will be entered for the Stream connection's username and password.

Create the connection

To create an OCI Streaming connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **OCI Streaming**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection Details** fields as follows:
 - a. Choose **Select a stream pool** to select an existing stream pool in the compartment. Click **Change Compartment** to select a stream pool in a different compartment.
 - b. Choose Enter stream pool information to manually enter bootstrap server connection details:
 - i. Under Bootstrap servers, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.

- **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - ii. Enter the Bootstrap server's Host and Port.
 - c. For **Username**, enter the Stream Pool username copied from the SASL Connection Settings in the prerequisite steps above.
 - d. For **Password**, enter the Auth token copied in prerequisite steps above.
 - e. (Optional) Click **Show advanced options** to drag and drop or select **Consumer** and **Producer properties**.
6. Click **Create**.

After the connection is created, it appears in the Connections list where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Connect to Oracle NoSQL

Learn to create a connection to Oracle NoSQL to use as a target in an OCI GoldenGate replication.

To create an Oracle NoSQL connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Technology Type**, select **Oracle NoSQL**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Tenancy**, select one of the following options:
 - Use current tenancy
 - Specify another tenancy, and then enter the **Tenancy OCID**

- b. For **Region**, select the OCI Object Storage region.
- c. For **User**, select one of the following options:
 - Use current user
 - Specify another user, and then enter the **User OCID**
- d. For **Private key file**, drag and drop or select the Privacy enhanced mail (PEM) file.

 **Tip:**

You can create and download the API key on your user details page, and then upload it here.

- e. Enter the **Private key passphrase**.
 - f. Enter the **Public key fingerprint**.
 - g. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

After the connection is created, it appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a Big Data deployment to use it in a replication.

Connect to Oracle Weblogic JMS (Java Message Service)

Learn to create a connection to Oracle Weblogic JMS to use as a source or target for an OCI GoldenGate replication process.

To create a Oracle Weblogic JMS connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Oracle WebLogic JMS**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:

- Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
- ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
 5. Complete the **Connection details** fields as follows:
 - a. Select whether to **Use Java Naming and Directory Interface (JNDI)**.
If selected, complete the following fields:
 - JNDI provider URL
 - Connection factory
 - Initial context factory
 - (Optional) JNDI security principal
 - (Optional) JNDI security credentials
 - b. Enter the **Connection URL**.
 - c. Enter the **JMS connection factory**.
 - d. For **Username**, enter the Oracle Weblogic JMS username.
 - e. For **Password**, enter the password.
 - f. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
 - g. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. Ensure you assign the connection to a Big Data deployment.

Connect to Amazon Kinesis

Learn to create a connection to Amazon Kinesis to use as a target in an OCI GoldenGate replication.

To create an Amazon Kinesis connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.

2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. From the **Type** dropdown, select **Amazon Kinesis**
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. On the Connection details page, complete the following fields:
 - a. for **Access key id**, enter the Amazon Web Services (AWS) user's Access key id.

 **Note:**

The Access key id must only contain alphanumeric characters and underscore and be 16-128 characters in length.

- b. For **Secret access key**, enter the AWS user's Secret access key.
 - c. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

It takes a few minutes for the connection to become Active. Ensure that you assign the connection to a Big Data deployment to use as a target in a data replication.

Connect to Amazon Redshift

Learn to create a connection to Amazon Redshift to use a target in an OCI GoldenGate replication.

To create an Amazon Redshift connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.

2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. From the **Type** dropdown, select **Amazon Redshift**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. On the Connection details page, complete the fields as follows:
 - a. for **Connection URL**, enter Amazon Redshift connection URL in the following format:

```
jdbc:redshift://[host]:[port]/[database]
```

 **Note:**

You can add additional parameters to the end of the connection URL as needed. See [Building the connection url](#) for more information.

- b. For **Username**, enter the Amazon Redshift username.
- c. For **Password**, enter the password for the Amazon Redshift username.
- d. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

It takes a few minutes for your connection to become Active. Ensure that you assign the connection to a Big Data deployment to use it as a target in a data replication.

Connect to Amazon S3

Learn to create a connection to Amazon S3 to use as a target for OCI GoldenGate.

To create an Amazon S3 connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Amazon S3**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Access key id**, enter the Amazon S3 access key ID.

 **Note:**

The Access key id must contain only alphanumeric characters and underscores, and be 16 to 128 characters in length.

- b. For **Secret access key**, enter the password for the Amazon S3 access key.
 - c. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. Ensure that you assign this connection to an OCI GoldenGate for Big Data deployment.

Connect to Apache Kafka

Learn to create a connection to Apache Kafka to use as an OCI GoldenGate source or target.

Create the connection

To create an Apache Kafka connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.

2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:

- a. For **Name**, enter a name for the connection.
- b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
- c. For **Compartment**, select the compartment in which to create the connection.
- d. For **Type**, select Apache Kafka.
- e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.

4. Click **Next**.

5. Complete the **Connection Details** fields as follows:

- a. Under Bootstrap servers:
 - Select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
If selected:
 - * Select the subnet
 - * (Optional) Enter the Private IP only if the hostname is not resolvable from your subnet or if it uses SSL/TLS.

 **Note:**

OCI GoldenGate rewrites the private IP in the format, `ip-10-0-0-0.ociggsvc.oracle.vcn.com`.

 **Tip:**

All nodes in the cluster must have FQDNs to allow for traversal over private endpoints.

- Enter the **Host** and **Port** number for the Bootstrap server. If Network connectivity via private endpoint is selected, then enter the **Private IP address** as well.
- (Optional) Click **+ Bootstrap server** to add another bootstrap server.
- b. For **Security protocol**, select one of the following and then complete the corresponding fields:
 - Plaintext
 - SASL over plaintext
 - SASL over SSL
 - SSL
- c. (Optional) Click **Show advanced options**.
 - If using this connection for Extract, then drag and drop or select the **Consumer** or **Producer properties** file under **Additional properties**.
 - To use Snappy compression in Kafka replication, drag and drop or select **Producer properties**, and change replication settings as discussed in [Using Compression OCI GoldenGate \(Confluent\) Kafka Replication](#).
 - To capture from Kafka, create a Kafka Consumer properties file with one of the following deserializers or converters:
 - Kafka Consumer properties for JSON deserializer:


```
key.deserializer=org.apache.kafka.common.serialization.ByteArrayDeserializer
value.deserializer=org.apache.kafka.common.serialization.ByteArrayDeserializer
```
 - Kafka Consumer properties for JSON converter:


```
key.converter=org.apache.kafka.connect.json.JsonConverter
value.converter=org.apache.kafka.connect.json.JsonConverter
```
 - Kafka Consumer properties for Avro converter:


```
key.converter=io.confluent.connect.avro.AvroConverter
value.converter=io.confluent.connect.avro.AvroConverter
```

6. Click **Create**.

After the connection is created, it appears in the Connections list where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Troubleshoot Kafka connection errors

Most connection issues result in `TimeoutException` errors. For example:

```
A failure occurred sending a message to Kafka to topic [ggstest]
org.apache.kafka.common.errors.TimeoutException: Topic ggstest not present
in metadata after 60000/120000 ms.
```

If you encounter this message in your Replicat report file, you can:

- Ensure the target topic is present or check that auto topic creation is enabled within the target Kafka settings.
- Ensure that there are no firewall rules blocking traffic.
- If you're running Kafka on OCI with a private endpoint, then ensure that you use the Internal FQDN as the bootstrap server in `server.properties` and in the Kafka connection.
- If you're connecting to a Confluent Cloud with private endpoints:
 - Ensure that the DNS zones and DNS records are configured properly in both OCI and the target third party cloud.
 - Ensure that the network connection between OCI and the target cloud work fine.
 - Test that you can connect to the target Confluent Cloud with OpenSSL (`openssl s_client -connect <bootstrap>`) from an OCI VM running in the same subnet connected to the third party cloud.
 - Test that you can publish or consume messages from a Kafka client running on OCI within the same subnet connected to the third party cloud. If it fails, then check your network settings on both OCI and the third party cloud.

Connect to Confluent Kafka

Learn to create a connection to Confluent Kafka, which you use with the Confluent Schema Registry connection, to serve as a source or target in an OCI GoldenGate Big Data deployment.

Create the connection

To create a Confluent Kafka connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.

- c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Confluent Kafka**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
 5. Complete the **Connection details** fields as follows:
 - Provide the Bootstrap details. You can find the bootstrap details of your cluster in the Confluent Kafka **Cluster settings**, and then **Endpoints**.
 - i. Select a Traffic routing method:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.

 **Note:**

Select this option if your Azure SQL Managed Instance doesn't have a public endpoint.

If selected:

- Select the subnet
- (Optional) Enter the Private IP only if the hostname is not resolvable from your subnet or if it uses SSL/TLS.

 **Note:**

OCI GoldenGate rewrites the private IP in the format, `ip-10-0-0-0.ociggsvc.oracle.vcn.com`.

 **Note:**

Learn more about Oracle GoldenGate connectivity.

- ii. Enter the Bootstrap server's **Host** and **Port**.

6. For **Security protocol**, select from:
 - Plaintext
 - SASL over plaintext, and then provide the Username and Password.
 - SASL over SSL, and then provide the Username, Password, and Truststore and Keystore values as needed.
 - SSL, and then provide the Truststore and Keystore values as needed.
7. (Optional) Click **Show advanced options**.
 - If using this connection for Extract, then drag and drop or select the **Consumer** or **Producer properties** file under **Additional properties**.
 - To use Snappy compression in Kafka replication, drag and drop or select **Producer properties**, and change replication settings as discussed in [Using Compression OCI GoldenGate \(Confluent\) Kafka Replication](#).
 - To capture from Kafka, create a Kafka Consumer properties file with one of the following deserializers or converters:
 - Kafka Consumer properties for JSON deserializer:

```
key.deserializer=org.apache.kafka.common.serialization.ByteArrayDeserializer
value.deserializer=org.apache.kafka.common.serialization.ByteArrayDeserializer
```
 - Kafka Consumer properties for JSON converter:

```
key.converter=org.apache.kafka.connect.json.JsonConverter
value.converter=org.apache.kafka.connect.json.JsonConverter
```
 - Kafka Consumer properties for Avro converter:

```
key.converter=io.confluent.connect.avro.AvroConverter
value.converter=io.confluent.connect.avro.AvroConverter
```

The connection appears in the Connections list, where you can select it to view its details. Ensure that you also create a connection to Confluent Schema Registry, and then assign both connections to a Big Data deployment.

Create a connection to Confluent Cloud with Private Links

Private Link lets you access your Confluent Cloud cluster running on a third party cloud through a private endpoint that exists in your virtual network.

Before you create the connection, ensure you have the following:

- Create private network connectivity between Oracle Cloud Infrastructure (OCI) and the target third party cloud.
- While adding network configuration for private link in Confluent Cloud, ensure that you select Private DNS Resolution.
- Configure DNS zones and set up DNS records in the third party cloud where you configured Confluent Cloud and in OCI. In OCI, you can create zones within your VCN's private views. Within zones, you can add the required DNS records.

You can use the instructions above to create the connection, but in place of Steps 5 and 6, do the following:

- For Step 5:
 - Provide the Bootstrap details. You can find the bootstrap details of your cluster in the Confluent Cloud cluster settings.
 - Select **Network connectivity via private endpoint**, and then select the subnet linked to the third party cloud where Confluent Cloud runs. Don't enter any private IP addresses.
- For Step 6:
 - For Security protocol, select **SASL over Plaintext**.
 - Enter the username and password.

Troubleshoot Kafka connection errors

Most connection issues result in `TimeoutException` errors. For example:

```
A failure occurred sending a message to Kafka to topic [ggstest]  
org.apache.kafka.common.errors.TimeoutException: Topic ggstest not  
present in metadata after 60000/120000 ms.
```

If you encounter this message in your Replicat report file, you can:

- Ensure the target topic is present or check that auto topic creation is enabled within the target Kafka settings.
- Ensure that there are no firewall rules blocking traffic.
- If you're running Kafka on OCI with a private endpoint, then ensure that you use the Internal FQDN as the bootstrap server in `server.properties` and in the Kafka connection.
- If you're connecting to a Confluent Cloud with private endpoints:
 - Ensure that the DNS zones and DNS records are configured properly in both OCI and the target third party cloud.
 - Ensure that the network connection between OCI and the target cloud work fine.
 - Test that you can connect to the target Confluent Cloud with OpenSSL (`openssl s_client -connect <bootstrap>`) from an OCI VM running in the same subnet connected to the third party cloud.
 - Test that you can publish or consume messages from a Kafka client running on OCI within the same subnet connected to the third party cloud. If it fails, then check your network settings on both OCI and the third party cloud.

Connect to Confluent Schema Registry

Learn to create a connection to Confluent Schema Registry that is used in conjunction with the Confluent Kafka connection to serve as a source or target for a OCI GoldenGate Big Data deployment.

To create a Confluent Schema Registry connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Confluent Schema Registry**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. Enter the Confluent Schema Registry **URL**.
 - b. Select an **Authentication type** from the dropdown. Select from:
 - None
 - Basic authentication (default)
 - Mutual authentication
 - c. For **Username**, enter the Confluent Schema Registry connection string, and then enter the **Password**.
 - d. (Optional) Drag-and-drop or select a **Truststore** file, and then enter the **Truststore password**.
 - e. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. Ensure that you create a connection to Confluent Kafka and then assign both the connections to a Big Data deployment to use as a source or target.

Connect to Azure Cosmos DB for MongoDB

Learn to create a connection to Azure Cosmos DB for MongoDB to use as a target for OCI GoldenGate.

To create an Azure Cosmos DB for MongoDB connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Technology Type**, select **Azure Cosmos DB for MongoDB**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection Details** fields as follows:
 - a. For **Connection string**, enter the database's connection string.
 - b. For **Username**, enter the username to connect to the database with.
 - c. For **Password**, enter the password associated to the username provided in the previous step.
 - d. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

After the connection is created, it appears in the Connections list where you can select it to view its details. Ensure that you assign the connection to a Big Data deployment to use it in a replication.

Connect to Azure Event Hubs

Learn to create a connection to Azure Event Hubs to use an OCI GoldenGate Big Data deployment source or target.

Create an Azure Event Hubs connection

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Azure Event Hubs**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. Under Bootstrap servers:
 - i. Select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.

 **Note:**

Select this option if your Azure SQL Managed Instance doesn't have a public endpoint.

If selected:

- Select the subnet
- (Optional) Enter the Private IP only if the hostname is not resolvable from your subnet or if it uses SSL/TLS.

 **Note:**

OCI GoldenGate rewrites the private IP in the format, `ip-10-0-0-0.ociggsvc.oracle.vcn.com`.

 **Note:**

Learn more about Oracle GoldenGate connectivity.

- ii. Enter the Azure Event Hubs **Host** and **Port**. If connecting via private endpoint, then enter the **Private IP address**. Click **+ Bootstrap server** to add another.
- b. For **Password**, enter the Azure Event Hubs namespace connection string.

 **Note:**

The Azure Event Hubs namespace connection string consists of three parts: Endpoint, SharedAccessKeyName, SharedAccessKey. For example, `Endpoint=sb://<NamespaceName>.servicebus.windows.net/;SharedAccessKeyName=<KeyName>;SharedAccessKey=<KeyValue>`.

- c. (Optional) Click **Show advanced options**.
 - If using this connection for Extract, then drag and drop or select the **Consumer** or **Producer properties** file under **Additional properties**.
 - To use Snappy compression in Kafka replication, drag and drop or select **Producer properties**, and change replication settings as discussed in [Using Compression OCI GoldenGate \(Confluent\) Kafka Replication](#).
 - To capture from Kafka, create a Kafka Consumer properties file with one of the following deserializers or converters:

- Kafka Consumer properties for JSON deserializer:

```
key.deserializer=org.apache.kafka.common.serialization.ByteArrayDeserializer
value.deserializer=org.apache.kafka.common.serialization.ByteArrayDeserializer
```

- Kafka Consumer properties for JSON converter:

```
key.converter=org.apache.kafka.connect.json.JsonConverter
value.converter=org.apache.kafka.connect.json.JsonConverter
```

- Kafka Consumer properties for Avro converter:

```
key.converter=io.confluent.connect.avro.AvroConverter
value.converter=io.confluent.connect.avro.AvroConverter
```

6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. You can then assign the connection to a Big Data deployment to use as a source or target.

Connect to Azure Data Lake Storage

Learn to create a connection to Azure Data Lake Storage to use as a target with OCI GoldenGate.

Create the Connection

To create an Azure Data Lake Storage connection

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.

2. On the Connections page, click **Create Connection**.

3. In the Create Connection panel, complete the **General Information** fields as follows:

- a. For **Name**, enter a name for the connection.
- b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
- c. For **Compartment**, select the compartment in which to create the connection.
- d. For **Technology Type**, select **Azure Data Lake Storage**.
- e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.

4. Click **Next**.

5. Complete the **Connection Details** fields as follows:

- a. For **Account name**, enter the Azure Cloud Storage Account Name.
- b. (Optional) For **Endpoint**, enter the Azure Storage service endpoint.
- c. For **Authentication type** select from the following:
 - **Shared key**:
 - **Storage Account Key**: Provide the Storage Account Access Key. See [Manage storage account access keys](#).
 - **Shared access signature**:
 - **SAS token**: Provide SAS token. Ensure that the SAS token is created on the account level. See [Create an account SAS](#).
 - **Azure Active Directory**:

 **Note:**

Before you configure the Azure Active Directory authentication type, ensure that you register an application in Azure AD App Registrations and assign the appropriate roles, for example "Storage BLOB Data Owner". See [Use the portal to create an Azure AD application and service principal that can access resources](#).

- **Azure tenant id**, located in Azure Active Directory/ App Registrations, select the application, and enter the Tenant ID.
- **Client id**, located in Azure Active Directory/ App Registrations and select the application.
- **Client secret**: Provide Azure Client Secret (Value) for the selected application id [Steps to generate a new client secret and link it to key-vault](#).

6. Click **Create**.

After the connection is created, it appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a target in a replication.

Troubleshoot connection issues

Most Azure Data Lake Storage connection issues happen because of Azure Data Lake Storage private endpoint configurations.

The following are common connectivity related error messages that you could encounter in the Replicat report file:

- =ERROR 2023-08-04 07:23:08.000008 [main] - Exception during initialisation of Azure blob service client for account[ociggtst]. com.azure.storage.blob.models.BlobStorageException: Status code 400, {"error":{"code":"InvalidUri","message":"The request URI is invalid.
- =ERROR 2023-08-01 20:23:24.000861 [main] - The Event Handler Framework failed to initialise.
- =ERROR 2023-08-04 08:13:30.000477 [main] - Exception during initialization of Azure blob service client for account[ociggtst]. com.azure.storage.blob.models.BlobStorageException:Status code 403, "<?xml version="1.0" encoding="utf-8"?><Error><Code>AuthorizationFailure</Code><Message>This request is not authorized to perform this operation.

If you're using Azure Data Lake Storage private endpoints and having issues with connection and/or replication, ensure that you:

- Check your OCI - Azure Interconnect details. Refer to [Step-by-Step Guide: Interconnecting Oracle Cloud Infrastructure and Microsoft Azure](#).

- Follow the steps outline in [OCI GoldenGate ADLS Connections with Private End Points](#)
- Configure your ADLS Private Endpoint Connection in Azure with target sub-resource BLOB. OCI GoldenGate only supports BLOB, so the connection fails if it is configured with dfs or other sub-resource types.

Connect to Azure Synapse Analytics

Learn to create a connection to Azure Synapse Analytics to use an OCI GoldenGate Big Data deployment target.

Create an Azure Synapse Analytics connection

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Technology type**, select **Azure Synapse Analytics**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. Enter the Azure Synapse Analytics **Connection string**.
 - b. Enter the Azure Synapse Analytics **Username** and **Password**.
 - c. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. You can then assign the connection to a Big Data deployment to use as a target.

Connect to MongoDB

Learn to create a connection to MongoDB to use as a source or target for OCI GoldenGate.

Create the connection

To create a MongoDB connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Technology Type**, select **MongoDB**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection Details** fields as follows:
 - a. For **Connection string**, enter a valid [MongoDB connection string URI](#) without the username and password. For example, `mongodb://mongodb1.example.com:27017`.
 - b. For **Username**, enter the username for this database.
 - c. For **Password**, enter the password associated with the username provided in the previous step.
 - d. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

After the connection is created, it appears in the Connections list where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Known issues

Issue with MongoDB Test connection

You may encounter an error when using Test connection with MongoDB connections. You can ignore this error and test MongoDB connections in the OCI GoldenGate deployment console. In the deployment console, open the navigation menu for the Administration Service, click **Configuration**. Your MongoDB connection should be listed as a credential, where you can click **Connect to <alias>** to test the connection.

Connect to Redis

Learn to create a connection to Redis to use as a target in an OCI GoldenGate replication.

To create a Redis connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. From the **Type** dropdown, select **Redis**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the Connection details as follows:
 - a. For **Servers**, enter a comma separated list of Redis server addresses, in <host>:<port> format.
 - b. From the **Authentication type** dropdown, select one of the following:
 - **None**
 - **Basic authentication**, if selected, enter the **Username** and **Password**.
 - c. For Security protocol, select one of the following:

- Plain
 - TLS
 - MTLS
- d. For **Network connectivity**, select a **Traffic routing method**:
- **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

The connection takes a few minutes to become Active. Ensure that you assign the connection to a Big Data deployment to use it as a target in a data replication.

Connect to Snowflake

Learn to create a connection to Snowflake to use as a target for OCI GoldenGate.

To create a Snowflake connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.

 **Note:**

Connecting to Snowflake Private Endpoints requires additional network configurations. For more information, see [Connecting OCI GoldenGate to Snowflake Private Endpoints](#).

2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Snowflake**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.

4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Connection URL**, enter the Snowflake database connection URL.
 - b. Select an **Authentication type** from the dropdown. Select:
 - Basic authentication, then enter the database username and password
 - Key pair authentication, then upload the private key and enter the private key password
 - c. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

After the connection is created, it appears in the Connections list where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Create a connection to Google Cloud Storage

Learn to create a connection to Google Cloud Storage to use as a target in an OCI GoldenGate replication.

To create a Google Cloud Storage connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. From the **Type** dropdown, select Google Cloud Storage.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.

5. On the Connection details page, complete the fields as follows:
 - a. For **Service account key file**, drag and drop the service account key file, or click **Select one** to upload it from your local machine.

 **Note:**

Learn more about [Service account keys](#).

- b. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 6. Click **Create**.

The connection takes a few minutes to become Active. Ensure that you assign the connection to a Big Data deployment to use it as a target in a data replication.

Connect to Google BigQuery

Learn to create a connection to Google BigQuery to use as a target in an OCI GoldenGate replication.

To create a Google Big Query connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. From the **Type** dropdown, select Google BigQuery.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.

5. On the Connection details page, complete the fields as follows:
 - a. For **Service account key file**, drag and drop the service account key file, or click **Select one** to upload it from your local machine.

 **Note:**

For more information, see [Authenticating with a service account key file](#).

- b. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

The connection takes a few minutes to become Active. Ensure that you assign the connection to a Big Data deployment to use it as a target in a data replication.

Connect to Elasticsearch Server

Learn to create a connection to Elasticsearch Server to use as a target in an OCI GoldenGate replication.

To create a Elasticsearch Server connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. From the **Type** dropdown, select **Elasticsearch Server**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the Connection details as follows:

- a. For **Servers**, enter a comma separated list of Elasticsearch Server server addresses, in <host>:<port> format.
 - b. From the **Authentication type** dropdown, select one of the following:
 - **None**
 - **Basic authentication**, if selected, enter the **Username** and **Password**.
 - c. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then enter the Fingerprint.
 - d. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

The connection takes a few minutes to become Active. Ensure that you assign the connection to a Big Data deployment to use it as a target in a data replication.

Connect to Hadoop Distributed File System

Learn to create a connection to an Hadoop Distributed File System to use as a target in an OCI GoldenGate replication process.

To create an Hadoop Distributed File System connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Technology Type**, select **Hadoop Distributed File System**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.

4. Click **Next**.
5. On the **Connection details** page, complete the fields as follows:
 - a. Drag and drop or select an XML configuration file for this database for **Hadoop configuration file**.
 - b. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

After the connection is created, it appears in the Connections list. Ensure that you assign the connection to a Big Data deployment to use it in a replication.

MySQL connections

Learn to create connections to MySQL databases, including:


- [Connect to OCI MySQL HeatWave](#)
- [Connect to MySQL Database Server](#)
- [Connect to Amazon Aurora MySQL](#)
- [Connect to Amazon RDS for MySQL](#)
- [Connect to Azure Database for MySQL](#)
- [Connect to Google Cloud SQL for MySQL](#)
- [Connect to Amazon RDS for MariaDB](#)
- [Connect to MariaDB](#)
- [Connect to SingleStoreDB](#)
- [Connect to SingleStoreDB Cloud](#)

Connect to OCI MySQL Heatwave

Learn to create a connection to OCI MySQL Heatwave to use as an OCI GoldenGate source or target.

To create an OCI MySQL Heatwave connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.

- c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **OCI OCI MySQL Heatwave**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
 5. Complete the **Connection details** fields as follows:
 - a. Choose **Select MySQL database system** to select an existing OCI MySQL Heatwave database system, and then enter the database name and a password for the `ggadmin` user.
 - b. Choose **Enter MySQL information** to manually enter the OCI MySQL Heatwave database system details:
 - i. For **Database name**, enter the OCI MySQL Heatwave database service name.
 - ii. Enter the OCI MySQL Heatwave database host (optional) and port in their respective fields.
-  **Note:**

For Host, enter the database's internal FQDN if available, any FQDN as an alias, or leave it blank and OCI GoldenGate assigns a system generated alias. [Learn more about FQDNs.](#)
- iii. For **Database username**, leave `ggadmin` as is.
 - iv. For **Database user password**, enter the password for the `ggadmin` user.
 - v. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
- c. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - d. (Optional) Under **Advanced options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Connect to MySQL Database Server

Learn to create a connection to MySQL Database Server to use as a OCI GoldenGate source or target.

To create a MySQL Database Server connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **MySQL Database Server**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection Details** fields as follows:
 - a. For **Database name**, enter the database name.

 **Note:**

Add double quotes around the username if it contains an @ symbol ("myuser@myserver", for example).

- b. Enter the database **Host** and **Port** number.
- c. For **Database username**, leave `ggadmin` as is.
- d. For **Database user password**, enter the password for `ggadmin`.
- e. For SSL details, select a **Security protocol** and **SSL mode**:
 - Plain
 - TLS, and then select the SSL mode

- MTLS, and then select the SSL mode
- f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
- 6. Click **Create**.

After the connection is created, it appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Connect to Amazon RDS for MySQL

Learn to create a connection to Amazon RDS for MySQL to use as an OCI GoldenGate source or target.

To create a Amazon RDS for MySQL connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Amazon RDS for MySQL**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the database name.

 **Note:**

Add double quotes around the username if it contains an @ symbol ("myuser@myserver", for example).

- b. Enter the database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the Amazon RDS for MySQL user.
 - d. For **Database user password**, enter the database user's password.
 - e. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - g. (Optional) Under **Advanced Options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Connect to Amazon RDS for MariaDB

Learn to create a connection to Amazon RDS for MariaDB to use as an OCI GoldenGate source or target.

To create a Amazon RDS for MariaDB connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Amazon RDS for MariaDB**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:

- Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
- ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
 5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the database name.

 **Note:**

Add double quotes around the username if it contains an @ symbol ("myuser@myserver", for example).

- b. Enter the database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the Amazon RDS for MariaDB user.
 - d. For **Database user password**, enter a password for the database user.
 - e. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - g. (Optional) Under **Advanced Options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a MySQL deployment to use it as a source or target in a replication.

Connect to Amazon Aurora MySQL

Learn to create a connection to Amazon Aurora MySQL to use as an OCI GoldenGate source or target.

To create a Amazon Aurora MySQL connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.

2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Amazon Aurora MySQL**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the Amazon Aurora MySQL database name.
 - b. Enter the Amazon Aurora MySQL database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the Amazon Aurora MySQL database username.

 **Note:**

Add double quotes around the username if it contains an @ symbol ("myuser@myserver", for example).

- d. For **Database user password**, enter the password for the Amazon Aurora MySQL database user.
 - e. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - g. (Optional) Under **Advanced options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Connect to Azure Database for MySQL

Learn to create a connection to Azure Database for MySQL to use as an OCI GoldenGate source or target.

To create a Azure Database for MySQL connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Azure Database for MySQL**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the database name.
 - b. Enter the database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the database username.

 **Note:**

Add double quotes around the username if it contains an @ symbol ("myuser@myserver", for example).

- d. For **Database user password**, enter the password for the database user.
- e. For **Security protocol**, select from the following options:
 - Plain

- TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - g. (Optional) Under **Advanced Options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Connect to Google Cloud SQL for MySQL

Learn to create a connection to Google Cloud SQL for MySQL to use as an OCI GoldenGate source or target.

To create a Google Cloud SQL for MySQL connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Google Cloud SQL for MySQL**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the Google Cloud SQL for MySQL database name.

 **Note:**

Add double quotes around the username if it contains an @ symbol ("myuser@myserver", for example).

- b. Enter the Google Cloud SQL for MySQL database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the Google Cloud SQL for MySQL database username.
 - d. For **Database user password**, enter the password for the Google Cloud SQL for MySQL database user.
 - e. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - g. (Optional) Under **Advanced options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. You can then assign this connection to a MySQL deployment.

Connect to MariaDB

Learn to create a connection to MariaDB to use as an OCI GoldenGate source or target.

To create a MariaDB connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.

- d. For **Type**, select **MariaDB**.
- e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the MariaDB database name.
 - b. Enter the MariaDB database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the MariaDB database username.

 **Note:**

Add double quotes around the username if it contains an @ symbol ("myuser@myserver", for example).

- d. For **Database user password**, enter the password for the MariaDB database user.
 - e. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - g. (Optional) Under **Show advanced options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Connect to SingleStoreDB

Learn to create a connection to SingleStoreDB to use as an OCI GoldenGate target.

To create a SingleStoreDB connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **SingleStoreDB**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the SingleStoreDB database name.
 - b. Enter the SingleStoreDB database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the SingleStoreDB database username.

 **Note:**

Add double quotes around the username if it contains an @ symbol ("myuser@myserver", for example).

- d. For **Database user password**, enter the password for the SingleStoreDB database user.
- e. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
- f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.

- g. (Optional) Under **Show advanced options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a target in a replication.

Connect to SingleStoreDB Cloud

Learn to create a connection to SingleStoreDB Cloud to use as an OCI GoldenGate source or target.

To create a SingleStoreDB connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **SingleStoreDB Cloud**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the SingleStoreDB database name.
 - b. Enter the SingleStoreDB Cloud database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the SingleStoreDB Cloud database username.

Note:

Add double quotes around the username if it contains an @ symbol ("myuser@myserver", for example).

- d. For **Database user password**, enter the password for the SingleStoreDB Cloud database user.

- e. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - g. (Optional) Under **Show advanced options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

PostgreSQL connections

Learn to create connections to PostgreSQL databases, including:

- [Connect to Amazon Aurora PostgreSQL](#)
- [Connect to Amazon RDS PostgreSQL](#)
- [Connect to Azure Database for PostgreSQL](#)
- [Connect to Google Cloud SQL for PostgreSQL](#)
- [Connect to PostgreSQL Server](#)

Create a connection to PostgreSQL Server

Learn to create a connection to PostgreSQL Server databases to use as a source or target with OCI GoldenGate.

To create a PostgreSQL Server connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.

- d. Select **PostgreSQL Server** from the **Type** dropdown.
- e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. On the Connection details page, complete the following fields:
 - a. For **Database name**, enter the PostgreSQL Server database name.
 - b. Enter the PostgreSQL Server database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the PostgreSQL Server database username.
 - d. For **Database user password**, enter the password for the PostgreSQL Server database user.
 - e. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - g. (Optional) Under **Show advanced options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. You can then assign the connection to a PostgreSQL deployment to use as a source or target.

Connect to Amazon RDS PostgreSQL

Learn to create a connection to an Amazon RDS PostgreSQL database that you can use as a source or target in a data replication.

To create a Amazon RDS PostgreSQL connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.

3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Amazon RDS PostgreSQL**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the Amazon RDS PostgreSQL database name.
 - b. Enter the Amazon RDS PostgreSQL database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the Amazon RDS PostgreSQL database username.
 - d. For **Database user password**, enter the password for the Amazon RDS PostgreSQL database user.
 - e. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - g. (Optional) Under **Show advanced options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. You can then assign this connection to a deployment.

Connect to Amazon Aurora PostgreSQL

Learn to create a connection to an Amazon Aurora PostgreSQL database to use as a source or target with OCI GoldenGate.

To create a Amazon Aurora PostgreSQL connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Amazon Aurora PostgreSQL**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the Amazon Aurora PostgreSQL database name.
 - b. Enter the Amazon Aurora PostgreSQL database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the Amazon Aurora PostgreSQL database username.
 - d. For **Database user password**, enter the password for the Amazon Aurora PostgreSQL database user.
 - e. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.

- **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - g. (Optional) Under **Advanced Options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. You can then assign this connection to a deployment.

Connect to Azure Database for PostgreSQL

To create a Azure Database for PostgreSQL connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Azure Database for PostgreSQL**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the Azure Database for PostgreSQL database name.
 - b. Enter the Azure Database for PostgreSQL database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the Azure Database for PostgreSQL database username.
 - d. For **Database user password**, enter the password for the Azure Database for PostgreSQL database user.
 - e. For **Security protocol**, select from the following options:

- Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
- f. For **Network connectivity**, select a **Traffic routing method**:
- **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
- g. (Optional) Under **Show advanced options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. You can then assign this connection to a deployment.

Connect to Google Cloud SQL for PostgreSQL

Learn to create a connection to Google Cloud SQL for PostgreSQL to use a source or target in a data replication.

To create a Google Cloud SQL for PostgreSQL connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the General Information fields, as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Google Cloud SQL for PostgreSQL**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the Google Cloud SQL for PostgreSQL database name.

- b. Enter the Google Cloud SQL for PostgreSQL database host and port in the **Host** and **Port** fields.
 - c. For **Database username**, enter the Google Cloud SQL for PostgreSQL database username.
 - d. For **Database user password**, enter the password for the Google Cloud SQL for PostgreSQL database user.
 - e. For **Security protocol**, select from the following options:
 - Plain
 - TLS, and then select the SSL mode
 - MTLS, and then select the SSL mode
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 - g. (Optional) Under **Show advanced options**, add **Connection attributes**. Click **+ Another attribute** to add more.
6. Click **Create**.

The connection appears in the Connections list, where you can select it to view its details. You can then assign the connection to a deployment to use as a source or target.

SQL Server connections

Learn to create connections to Microsoft SQL Server databases to use as sources or targets in OCI GoldenGate replications.

- [Connect to Amazon RDS for SQL Server](#)
- [Connect to Azure SQL Database](#)
- [Connect to Azure SQL Managed Instance](#)
- [Connect to Microsoft SQL Server](#)
- [Connect to Google Cloud SQL for SQL Server](#)

Connect to Azure SQL Database

Learn to create a connection to Azure SQL Database to use as a OCI GoldenGate source or target.

Ensure that you prepare and configure the system for Oracle GoldenGate.

To create an Azure SQL Database database connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.

2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Azure SQL Database**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection Details** fields as follows:
 - a. For **Database name**, enter the database name.
 - b. Enter the database **Host** and **Port** number.
 - c. For **Username**, enter the Azure SQL Database username.
 - d. For **Password**, enter the password for the Azure SQL Database user.
 - e. For SSL details, select a **Security protocol** and **SSL mode**:
 - Plain
 - TLS, and then select whether to validate the server certificate
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

After the connection is created, it appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Create a connection to Azure SQL Managed Instance

Learn to create a connection to Azure SQL Managed Instance to use as a OCI GoldenGate source or target.

Create the connection

To create an Azure SQL Managed Instance connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Azure SQL Managed Instance**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the database schema name.
 - b. Enter the database **Host** and **Port** number.
 - c. For **Database username**, enter the Azure SQL Managed Instance database username.
 - d. For **Database user password**, enter the password for the Azure SQL Managed Instance database.
 - e. For SSL details, select a **Security protocol** and **SSL mode**:
 - Plain
 - TLS, and then select whether to validate the server certificate
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.

 **Note:**

Select this option if your Azure SQL Managed Instance doesn't have a public endpoint.

If selected:

- Select the subnet
- (Optional) Enter the Private IP only if the hostname is not resolvable from your subnet or if it uses SSL/TLS.

 **Note:**

OCI GoldenGate rewrites the private IP in the format,
`ip-10-0-0-0.ociggsvc.oracle.vcn.com`.

 **Note:**

Learn more about Oracle GoldenGate connectivity.

6. Click **Create**.

After the connection is created, it appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Known issues

Private endpoints not supported for Azure SQL Managed Instance

Although Network connectivity settings appear on the Create connection screens for Azure SQL Managed Instance, private endpoints for Azure SQL Managed Instance are not currently supported.

Workaround: None.

Connect to Microsoft SQL Server

Learn to create a connection to Microsoft SQL Server to use as a OCI GoldenGate source or target.

Ensure that you prepare and configure the system for Oracle GoldenGate.

To create a Microsoft SQL Server connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Microsoft SQL Server**.

- e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the database name.
 - b. Enter the database **Host** and **Port** number.
 - c. For **Username**, enter the Microsoft SQL Server database username.
 - d. For **Password**, enter the password for the Microsoft SQL Server database user.
 - e. For SSL details, select a **Security protocol** and **SSL mode**:
 - Plain
 - TLS, and then select whether to validate the server certificate
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

After the connection is created, it appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Connect to Amazon RDS for SQL Server

Learn to create a connection to Amazon RDS for SQL Server to use as an OCI GoldenGate source or target.

Ensure that you prepare and configure the system for Oracle GoldenGate.

To create an Amazon RDS for SQL Server connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.

- b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Amazon RDS for SQL Server**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
 5. Complete the **Connection Details** fields as follows:
 - a. For **Database name**, enter the database name.
 - b. Enter the database **Host** and **Port** number.
 - c. For **Username**, enter the Amazon RDS for SQL Server database username.
 - d. For **Password**, enter the password for the Amazon RDS for SQL Server database user.
 - e. For SSL details, select a **Security protocol** and **SSL mode**:
 - Plain
 - TLS, and then select whether to validate the server certificate
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
 6. Click **Create**.

After the connection is created, it appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Connect to Google Cloud SQL for SQL Server

Learn to create a connection to Google Cloud SQL for SQL Server to use as a OCI GoldenGate source or target.

Ensure that you prepare and configure the system for Oracle GoldenGate.

To create a Google Cloud SQL for SQL Server connection:

1. From the OCI GoldenGate Overview page, click **Connections**.

You can also click **Create Connection** under the Get started section and skip to step 3.

2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. For **Type**, select **Google Cloud SQL for SQL Server**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the **Connection details** fields as follows:
 - a. For **Database name**, enter the database name.
 - b. Enter the database **Host** and **Port** number.
 - c. For **Username**, enter the Google Cloud SQL for SQL Server database username.
 - d. For **Password**, enter the password for the Google Cloud SQL for SQL Server database user.
 - e. For SSL details, select a **Security protocol** and **SSL mode**:
 - Plain
 - TLS, and then select whether to validate the server certificate
 - f. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.
6. Click **Create**.

After the connection is created, it appears in the Connections list, where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

Create a Generic connection

Use a Generic connection for the Data Transforms deployment to use as a data source.

See supported connection types for Data Transforms to ensure your data source is supported.

To create a Generic connection:

1. From the OCI GoldenGate Overview page, click **Connections**.
You can also click **Create Connection** under the Get started section and skip to step 3.
2. On the Connections page, click **Create Connection**.
3. In the Create Connection panel, complete the **General Information** fields as follows:
 - a. For **Name**, enter a name for the connection.
 - b. (Optional) For **Description**, enter a description that helps you distinguish this connection from others.
 - c. For **Compartment**, select the compartment in which to create the connection.
 - d. From the **Type** dropdown, select **Generic connection**.
 - e. (Optional) Click **Show advanced options** to manage keys or add tags.
 - i. Under **Security**, select one of the following:
 - Select **Use Oracle-managed encryption key** to leave all encryption key management to Oracle.
 - Select **Use customer-managed encryption key** to select a specific encryption key stored in your OCI Vault to encrypt your connection credentials.
 - ii. Under **Tags**, add tags to organize your resources.
4. Click **Next**.
5. Complete the Connection details page as follows:
 - a. For **Host**, enter a comma separated list of <host>:<port> entries.
 - b. For **Network connectivity**, select a **Traffic routing method**:
 - **Shared endpoint**, to share an endpoint with the assigned deployment. You must allow connectivity from the deployment's ingress IP.
 - **Dedicated endpoint**, for network traffic through a dedicated endpoint in the assigned subnet in your VCN. You must allow connectivity from this connection's ingress IPs.

 **Note:**

Select this option if your Azure SQL Managed Instance doesn't have a public endpoint.

If selected:

- Select the subnet

- (Optional) Enter the Private IP only if the hostname is not resolvable from your subnet or if it uses SSL/TLS.

 **Note:**

OCI GoldenGate rewrites the private IP in the format, `ip-10-0-0-0.ociggsvc.oracle.vcn.com`.

 **Note:**

Learn more about Oracle GoldenGate connectivity.

6. Click **Create**.

It takes a few minutes for the connection to become Active. Ensure that you assign the connection to a deployment to use it in a data transform. After assigning the connection to the Data Transforms deployment, launch the console, and then log in. You must then create a connection for each data source in the Data Transforms console using the Host name(s) provided in the Generic connection(s).

Connect to Oracle GoldenGate deployments

Learn to create connection to Oracle GoldenGate deployments to use Distribution or Receiver paths as OCI GoldenGate sources.

About Oracle GoldenGate connections

The Oracle GoldenGate connection type lets you create connections to other Oracle GoldenGate deployments. For example, replicate data between a MySQL database and Kafka, you need a MySQL deployment type and a Big Data deployment type. The Oracle GoldenGate deployment lets you create a connection between the two deployment types. These two deployment types need not be in the same compartment or tenancy. Create the connection to the Oracle GoldenGate deployment that initiates the replication.

Create a connection to GoldenGate

To create a connection to a GoldenGate server:

1. On the Connections page, click **Create Connection**.
2. In the Create connection panel, for General information, complete the following fields, and then click **Next**.
 - a. For **Name**, enter a name for this connection.
 - b. For **Description**, enter a brief, friendly description for this connection.
 - c. For **Compartment**, select the compartment in which to create this connection.
 - d. For **Type**, select **GoldenGate**.
3. For Connection details:

- a. Choose **Select GoldenGate deployment** to choose an existing GoldenGate deployment in a specific compartment. Click **Change compartment** to select a GoldenGate deployment in a different compartment.
 - b. Choose **Enter GoldenGate information** to manually enter the connection details for the GoldenGate deployment you want to connect to.
4. Under Network connectivity, select **Network connectivity via private endpoint** if the GoldenGate deployment can only be accessed through a private IP.
 - a. If selected, then for Subnet in <compartment-name>, select the subnet to which a private endpoint is created from the OCI GoldenGate service tenancy. This creates a network route for the OCI GoldenGate deployment to connect to the database within your customer tenancy. Click **Change Compartment** to select a subnet in a different compartment.
 - b. For **Private IP address**, enter the private IP for the GoldenGate deployment.
5. Click **Create**.

The connection appears in the Connections list where you can select it to view its details. Ensure that you assign the connection to a deployment to use it as a source or target in a replication.

5

Replicate data

Learn to create and run GoldenGate processes.

Articles in this section:

- [Create data replication resources](#)
- [Explore the OCI GoldenGate deployment console](#)
- [Add Extracts](#)
- [Add a Distribution Path](#)
- [Add a Receiver Path](#)
- [Add Replicats](#)
- [Using the Admin Client](#)

Create data replication resources

Learn to create data replication resources, such as deployments and connections, essential to getting started with Oracle Cloud Infrastructure GoldenGate.

Create a deployment

Deployments let you create and organize Oracle Cloud Infrastructure GoldenGate resources, and enables the OCI GoldenGate deployment console, where you can add and manage data replication processes. You can create deployments for Oracle Database, Big Data, Microsoft SQL Server, MySQL, and PostgreSQL sources or targets.

Note:

In Oracle Cloud Infrastructure Identity and Access Management (OCI IAM) enabled tenancies, step 16 requires you to select the credential store you'll use to log in to the deployment console.

- If you intend to select OCI IAM, ensure that you configure identity domains for OCI GoldenGate.
- If you intend to select GoldenGate, ensure that you first [set up your Vault](#). [Learn more about Vault service](#).

Note:

A virtual private vault is not required.

Depending on your source and target technology types, you may need more than one deployment. If you're not sure how many deployments you need for your solution, see Example topologies for details.

To create a deployment:

1. In the Console navigation menu, click **Oracle Database**, and then select **GoldenGate**.
2. On the Deployments page, click **Create deployment**.
3. In the Create deployment panel, enter a name and optionally, a description.
4. From the Compartment dropdown, select a compartment in which to create the deployment.
5. Select one of the following options:
 - **Production:** Sets up a deployment with recommended defaults for a production environment. The minimum number of OCPUs is 4, with auto-scaling enabled.
 - **Development or testing:** Sets up a deployment with recommended defaults for a development or testing environment. The minimum number of OCPUs is 1.
6. For **OCPU count** enter the number of Oracle Compute units (OCPUs) to use.

 **Note:**

One OCPU is equivalent to 16gb of memory. For more information, see OCPU management and billing.

7. (Optional) Select **Auto scaling**.

 **Note:**

Auto scaling enables OCI GoldenGate to scale up to three times the number of OCPUs you specify for OCPU Count, up to 24 OCPUs. For example, if you specify your OCPU Count as 2 and enable Auto Scaling, then your deployment can scale up to 6 OCPUs. If you specify your OCPU Count as 20 and enable Auto Scaling, OCI GoldenGate can only scale up to 24 OCPUs.

8. From the **Subnet in <Compartment>** dropdown, select the subnet to which a private endpoint is created from the OCI GoldenGate service tenancy. This ensures that the deployment is always available over this subnet, as long as the policies for this subnet allow access. The private endpoint is only used to access the deployment console, and doesn't provide access to other resources in the subnet.

To select a subnet in a different compartment, click **Change compartment**.

 **Note:**

You can only select a private subnet when creating a deployment.

9. Select a license type.
10. (Optional) Click **Show advanced options** for network options and to add tags.
 - a. In the **Network** tab,
 - i. Select **Enable GoldenGate console public access** to include a public endpoint in addition to a private endpoint, and allow public access to the deployment console for users. If selected, OCI GoldenGate creates a load balancer in your tenancy to create a public IP. Select a subnet in the same VCN as this deployment in which to create the load balancer.

 **Note:**

The load balancer is a resource that comes with an additional cost. You can manage this resource, but ensure that you don't delete the load balancer while your deployment is still in use. [Learn more about load balancer pricing.](#)

- ii. Select **Customize endpoint** to provide a private fully qualified domain name (FQDN) prefix that you'll use to access the private service console URL. You can also optionally upload an SSL/TLS certificate (.pem) and its corresponding private key, however, password protected certificates are not supported. It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.
A self-signed certificate is generated for you, if you don't provide one.

 **Note:**

It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.

- b. In the **Maintenance** tab:
 - i. Select **Customize maintenance window** to define the start of the maintenance window to upgrade the deployment.
 - ii. (Optional) For **Major release auto-upgrade period in days**, enter the number of days, between 0 and 365.
 - iii. (Optional) For **Bundle release auto-upgrade period in days**, enter the number of days, between 0 and 180 days.
 - iv. (Optional) For **Security patch auto-upgrade period in days**, enter the number of days, between 0 and 14 days.
 - v. Select **Enable interim release auto-upgrade**, and, optionally, enter the number of days.

 **Note:**

[Learn more about scheduling upgrades.](#)

- c. In the **Tags** tab, add tags to help track the resources within your tenancy. Click **+ Additional tag** to add more tags. [Learn more about tagging.](#)

11. Click **Next**.
12. For Deployment type, select **Data replication**.
13. From the **Select a technology** dropdown, select one of the following technology types:
 - Oracle Database
 - Big Data
 - MySQL
 - PostgreSQL
 - Microsoft SQL Server

See what's supported to learn which databases and technologies you can use as OCI GoldenGate sources and targets.

14. For **Version**, the latest version is automatically selected. Click **Change version** to select a different version.

 **Note:**

Learn more about versions.

15. For **GoldenGate instance name**, enter the name that the deployment will assign to the GoldenGate deployment instance upon creation.
16. For Credential store, select one of the following:
 - **OCI Identity and Access Management (OCI IAM)**, to enable users to log in to the the deployment console using their Oracle Cloud account (single sign on) in IAM (Identity and Access Management) enabled tenancies.

 **Note:**

Once you select IAM, you won't be able to switch to GoldenGate when you edit the deployment settings at a later time.

- **GoldenGate**, for GoldenGate to manage users.
 - a. Enter the **Administrator username**
 - b. Select a password secret in your compartment or click **Change compartment** to select one in a different compartment. You can also create a new password secret.
To create a new password secret:
 - i. Click **Create password secret**.
 - ii. In the Create secret panel, enter a name for the secret, and optionally, a description.
 - iii. Select a compartment from the **Compartment** dropdown in which to save your secret.
 - iv. Select a vault in the current compartment, or click **Change compartment** to select a vault in a different compartment.

- v. Select an **Encryption key**.

 **Note:**

Only AES keys, Software protected keys, and HSM keys are supported. RSA and ECDSA keys are not supported for GoldenGate password secret keys.

- vi. Enter a password 8 to 30 characters in length, containing at least 1 uppercase, 1 lowercase, 1 numeric and 1 special character. The special characters must not be '\$', '^' or '?'.
- vii. Confirm the password.
- viii. Click **Create**.

 **Note:**

You can manage GoldenGate users in the deployment console. Learn more.

- 17. Click **Create**.

After the deployment is created and becomes Active, it starts automatically. You can then select **Launch console** in the deployment's Actions (three dots) menu on the Deployments page, or click **Launch console** on the deployment details page to access the OCI GoldenGate deployment console.

Ensure that you create and assign connections to use with your deployment.

Assign a connection to a deployment

Ensure that you have connections created for your source and target technologies.

To assign a connection to a deployment:

1. On the deployment details page, under **Resources**, click **Assigned connections**.
2. Click **Assign connection**.
3. In the Assign connection dialog, select a connection from the dropdown. If you want to select a connection from a different compartment, click **Change Compartment**.
4. Click **Assign connection**.

The selected connection appears in the Assigned connections list. You can also view and manage this relationship from the Connection details page under **Assigned deployments**.

Access the deployment

After you create the Data replication deployment and assign connections, you can access the OCI GoldenGate deployment console from the deployment details page.

To access the OCI GoldenGate deployment:

1. On the OCI GoldenGate Deployments page, select the Data replication deployment.
2. On the deployment details page, click **Launch console**.

Alternatively, you can copy the **Console url** and paste it into your browser.

- For an IAM-enabled deployment, you're prompted to give the application access to `get_groups`. Click **Allow** to continue.
- For a non-IAM enabled deployment, on the log in page, enter the Administrator username and Administrator password provided when you created the deployment.

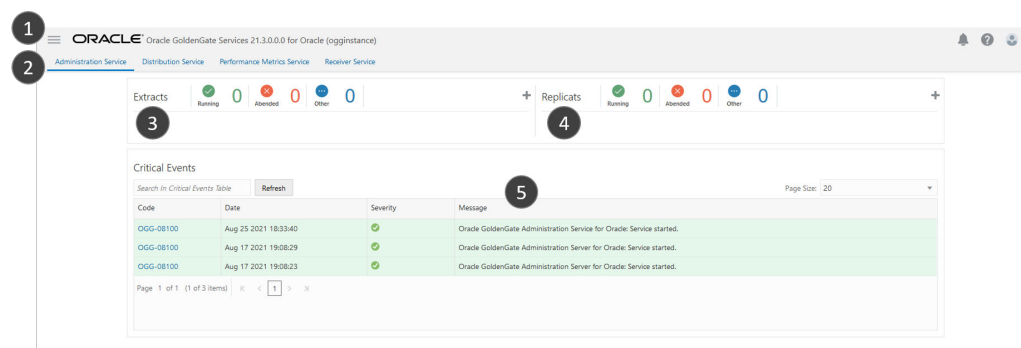
You can now explore the deployment console or start adding Extracts, Replicats, and Paths to build your data replication solution.

Explore the OCI GoldenGate deployment console

If you're new to GoldenGate Microservices, take a moment to get to know the OCI GoldenGate deployment console and where you'll create and manage your replication processes.

You launch the OCI GoldenGate deployment console from a deployment's Actions (ellipsis icon) menu on the Deployments page, or from the deployment's details page. Log in with the Oracle Cloud Infrastructure GoldenGate username and password that you specified when you created the deployment.

After you log in successfully, you're brought to the OCI GoldenGate deployment console Administration Service Overview.



The main areas of the Administration Service Console are:

- Navigation menu:** Click the navigation icon to show or hide the navigation menu.
- The navigation bar lets you switch between the following services:
 - Administration Service:** (currently shown) Administers, manages, and monitors Extract and Replicat processes within an Oracle GoldenGate deployment.
 - Distribution Service:** A networked data distribution agent that conveys and processes data and commands in a distributed deployment environment.
 - Performance Metrics Service:** Collects and stores deployment performance results.
 - Receiver Service:** Interoperates with the Distribution Server to handle all incoming trail files.
- Extracts:** Displays the number of running, failed, or other Extracts. Click Add Extract (plus icon) to create a new Extract.

4. **Replicats:** Displays the number of running, failed, or other Replicats. Click Add Replicat (plus icon) to create a new Replicat.
5. **Critical Events:** Displays the severity of critical events.

 **WARNING:**

You must only create and edit connections in the Oracle Cloud console. Refrain from creating or editing connections in the Credentials screen of the deployment console. Updates are automatically synced to the deployment from the Oracle Cloud console.

Add Extracts

Learn to configure Extracts for different types of supported connections.

Relational Database Management System (RDBMS) Extracts:

- [Add an Extract for Oracle Database](#)
- [Add an Extract for Microsoft SQL Server](#)
- [Add an Extract for MySQL](#)
- [Add an Extract for PostgreSQL](#)

Big Data Extracts:

- [Add an Extract for Kafka](#)
- [Add an Extract for MongoDB](#)

RDBMS Extracts

Learn to create Extracts for relational database management systems.

Add an Extract for Oracle Database

Extract is a process that runs against the source data source connection and extracts, or captures, data. Learn to add an Extract for Oracle Database, OCI Autonomous Databases, Oracle Exadata, and Amazon RDS for Oracle technologies.

Before you begin

Before adding and running an Extract to capture data from the source, ensure that you complete the following:

- Oracle GoldenGate relies on the redo logs to capture the data that it needs to replicate source transactions. Enable supplemental logging on the source database for unidirectional replication, or both source and target for bidirectional replication:

```
ALTER DATABASE ADD SUPPLEMENTAL LOG DATA
```

- Prepare the database for Oracle GoldenGate.

Adding transaction information

Enable table-level supplemental by adding TRANDATA. If schema-level supplemental logging is already enabled, you can skip this task.

For more information see, Configure logging properties.

To add TRANDATA:

1. Log in to the GoldenGate Deployment Console if you aren't already logged in.
2. In the navigation menu, click **Configuration**.
3. In the **Administration Service** tab, go to the **Database** tab, and then click the connect icon for the source database.
4. Next to TRANDATA Information, click **Add TRANDATA** (plus icon).
5. For **Schema Name**, enter the database schema name, and then click **Submit**.

Add an Extract

To add an Extract for Oracle Database:

1. In the OCI GoldenGate deployment console, ensure that you're on the Administration Service Overview page, and then click **Add Extract** (plus icon).
2. On the Add Extract page, for **Extract Type**, select one of the following, and then click **Next**.
 - Integrated Extract
 - Initial Load Extract
3. On the Extract Options page, under Basic Information, complete the fields as needed:
 - a. For **Process Name**, enter a name for the Extract process, up to 8 characters.
 - b. (Optional) For **Description**, enter a brief description to help you distinguish this process from others.
 - c. For **Intent**, select the option that best describes the purpose of this Extract:
 - Unidirectional (default)
 - High Availability
 - Disaster Recovery
 - N-Way
 - d. For **Begin**, select the location in the redo or transaction log where the Extract starts to capture data:
 - Now
 - Custom time
 - CSN
 - e. For **Trail Name**, enter a two character name for the Trail file.
 - f. (Optional) For **Trail Subdirectory**, set a custom location for the generated Trail file.

- g. (Optional) For **Trail Size**, set the max size for the generated trail file.
 - h. For **Remote**, enable this option if the Extract trail is to be written directly to a remote Oracle GoldenGate installation.
 4. Under Source Database Credential, you can either create a new credential or select an existing Credential Domain and Alias for the source database.
 5. (Optional) Under Registration Information, complete the fields as needed:
 - a. For **CSN**, enter the Commit Sequence Number (CSN).
 - b. For **Share**, choose the method to share the LogMiner data dictionary:
 - **Automatic**: allows the system to choose the method for sharing.
 - **None**: doesn't share the dictionary.
 - **Extract**: shares the LogMiner dictionary for this Extract.
 - c. For **Optimized**, enable this option to optimize Extract registration.
 - d. For **Downstream Capture**, enable this option to set up a downstream Extract for log mining.
 6. Under Downstream Mining, complete the following fields as needed:
 - a. For **Mining Credential Domain**, enter the downstream mining database's domain name.
 - b. For **Mining Credential Alias**, enter the downstream mining database's alias
 - c. For **No UserID**, enable this option if there is no source database connection. If selected, then the ADG fetch options are enabled.
 - d. For **ADG Fetch Credential Domain**, enter the ADG fetch database's domain name.
 - e. For **ADG Fetch Credential Alias**, enter the ADG fetch database's alias.
 7. (Optional) Under Encryption Profile, enter the encryption profile description. The Local Wallet profile is selected by default if an encryption profile wasn't created.
 - a. Select the profile name from the dropdown. You can select the Local Wallet or a custom profile.
 - b. Select the encryption profile from the dropdown.
 - c. Specify the masterkey for the encryption profile.
 8. Under Managed Options, enable **Critical to deployment health** to view Metrics on the Deployment Details page and Monitoring dashboard in the Oracle Cloud console. Complete the other optional fields as needed.

 **Note:**

Adding a profile and configuring Auto Start and Auto Restart options enables your deployment to restart automatically after a network disruption. See [Configure managed processes](#) to learn more.

9. Click **Next**.
10. On the Extract Parameters page, you can edit the parameter file in the text area to list the table details to capture. For example:

```
table source.table1;
```

 **Note:**

GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

11. Select **Register Extract in the background** to register the Extract in the background asynchronously.
12. Click **Create and Run** to create and start the Extract. If you click **Create**, then you can manually start Extract later from the Administration Service Overview page.

You're returned to the Administration Service Overview page, where you can view the status of the Extract process. Select **Details** from the Extract **Action** menu to view process information, checkpoint, statistics, parameters, and reports.

Learn more

Interested in learning more about the Extract process and capturing data? Refer to the following resources:

- Replicate data between cloud databases in the same region
- Send data from Oracle GoldenGate to OCI GoldenGate
- Configure bidirectional replication

Add an Extract for Microsoft SQL Server

Learn to add an Extract for Microsoft SQL Server, Amazon RDS for SQL Server, Azure SQL Database, or Azure SQL Managed Instance technologies.

To add an extract for Microsoft SQL Server database:

1. In the OCI GoldenGate deployment console, ensure that you're on the Administration Service Overview page, and then click **Add Extract** (plus icon).
2. On the Add Extract page, for **Extract Type**, select one of the following, and then click **Next**.
 - Change Data Capture
 - Initial Load
3. On the Extract Options page, under Basic Information, complete the fields as needed:
 - a. For **Process Name**, enter a name for the Extract process, up to 8 characters.
 - b. (Optional) For **Description**, enter a brief description to help you distinguish this process from others.
 - c. For **Credential Domain**, select **Oracle GoldenGate**.
 - d. For **Intent**, select the option that best describes the purpose of this Extract:

- Unidirectional (default)
 - Disaster Recovery
 - N-Way
- e. For **Credential Alias**, select your source Microsoft SQL Server connection.
 - f. For **Trail Name**, enter a two character name for the Trail file.
 - g. (Optional) For **Trail Subdirectory**, set a custom location for the generated Trail file.
 - h. (Optional) For **Trail Size**, set the max size for the generated trail file.
4. (Optional) Under Encryption Profile, enter the encryption profile description. The Local Wallet profile is selected by default if an encryption profile wasn't created.
 - a. Select the profile name from the dropdown. You can select the Local Wallet or a custom profile.
 - b. Select the encryption profile from the dropdown.
 - c. Specify the masterkey for the encryption profile.
 5. (Optional) Under Managed Options, you can configure the following:
 - Profile Name
 - Auto Start
 - Auto Restart

 **Note:**

Adding a profile and configuring Auto Start and Auto Restart options enables your deployment to restart automatically after a network disruption. See [Configure managed processes](#) for more information.

6. Click **Next**.
7. On the Extract Parameters page, you can edit the parameter file in the text area to list the table details to capture. For example:

```
table source.table1;
```

 **Note:**

GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")  
setenv (TZ="GMT+5")
```

[Learn more about SETENV.](#)

8. Select **Register Extract in the background** to register the Extract in the background asynchronously.

9. Click **Create and Run** to create and start the Extract. If you click **Create**, then you can manually start Extract later from the Administration Service Overview page.

You're returned to the Administration Service Overview page, where you can view the status of the Extract process. Select **Details** from the Extract **Action** menu to view process information, checkpoint, statistics, parameters, and reports.

Add an Extract for MySQL

Extract is a process that runs against the source data source connection and extracts, or captures, data. Learn to add an Extract for MySQL Database Server, OCI MySQL Heatwave, Amazon Aurora MySQL, Amazon RDS for MySQL, Amazon RDS for MariaDB, Azure Database for MySQL, and Google Cloud SQL for MySQL technologies.

Before you begin

Before adding and running an Extract, ensure that you prepare and configure the system for Oracle GoldenGate.

Add an Extract

To add an extract for MySQL database:

1. In the OCI GoldenGate deployment console, ensure that you're on the Administration Service Overview page, and then click **Add Extract** (plus icon).
2. On the Add Extract page, for **Extract Type**, select one of the following, and then click **Next**.
 - Change Data Capture
 - Initial Load
3. On the Extract Options page, under Basic Information, complete the fields as needed:
 - a. For **Process Name**, enter a name for the Extract process, up to 8 characters.
 - b. (Optional) For **Description**, enter a brief description to help you distinguish this process from others.
 - c. For **Intent**, select the option that best describes the purpose of this Extract:
 - Unidirectional (default)
 - Disaster Recovery
 - N-Way
 - d. Enable **Remote** only if capturing data from a MySQL database that doesn't use global transaction identifiers (GTIDs).
 - e. For **Credential Domain**, select **Oracle GoldenGate**.
 - f. For **Credential Alias**, select your source MySQL connection.
 - g. For **Trail Name**, enter a two character name for the Trail file.
 - h. (Optional) For **Trail Subdirectory**, set a custom location for the generated Trail file.

- i. (Optional) For **Trail Size**, set the max size for the generated trail file.
4. (Optional) Under Encryption Profile, enter the encryption profile description. The Local Wallet profile is selected by default if an encryption profile wasn't created.
 - a. Select the profile name from the dropdown. You can select the Local Wallet or a custom profile.
 - b. Select the encryption profile from the dropdown.
 - c. Specify the masterkey for the encryption profile.
5. (Optional) Under Managed Options, you can configure the following:
 - Profile Name
 - Auto Start
 - Auto Restart

 **Note:**

Adding a profile and configuring Auto Start and Auto Restart options enables your deployment to restart automatically after a network disruption. See [Configure managed processes](#) for more information.

6. Click **Next**.
7. On the Extract Parameters page, you can edit the parameter file in the text area to list the table details to capture. For example:

```
table source.table1;
```

 **Note:**

GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")  
setenv (TZ="GMT+5")
```

[Learn more about SETENV.](#)

8. Select **Register Extract in the background** to register the Extract in the background asynchronously.
9. Click **Create and Run** to create and start the Extract. If you click **Create**, then you can manually start Extract later from the Administration Service Overview page.

You're returned to the Administration Service Overview page, where you can view the status of the Extract process. Select **Details** from the Extract **Action** menu to view process information, checkpoint, statistics, parameters, and reports.

Learn more

Interested in learning more about the Extract process and capturing data? See:

- Replicate data from MySQL to Autonomous Data Warehouse

Known Issues

Remote change data capture Extracts fail for GTID enabled databases

When you create a Change Data Capture Extract process with the Remote option enabled for a MySQL database that uses global transaction identifiers (GTIDs), the Extract process fails and the following error is reported:

```
ERROR   OGG-25192  Trail file '<trail name>' is remote. Only local
trail allowed for this extract.
```

Workaround: On the Parameter file screen of the Change Data Capture Extract, remove the line, `TRANLOGOPTIONS ALTLOGDEST REMOTE`.

For more information, see [Using Oracle GoldenGate for MySQL](#).

Add an Extract for PostgreSQL

Extract is a process that runs against the source data source connection and extracts, or captures, data. Learn how to add an Extract for PostgreSQL Database, Amazon Aurora PostgreSQL, Amazon RDS PostgreSQL, Azure Database for PostgreSQL, and Google Cloud SQL for PostgreSQL technologies.

Before you begin

Before adding and running an Extract to capture data from the source, ensure that you:

- Prepare your database for Oracle GoldenGate
- Created a PostgreSQL connection and associated the connection to the PostgreSQL deployment
- Enable supplemental logging:
 1. Launch the PostgreSQL GoldenGate deployment console:
 - a. From the Deployments page, select the PostgreSQL deployment to view its details.
 - b. On the PostgreSQL deployment details page, click **Launch console**.
 - c. On the deployment console sign in page, enter the GoldenGate admin credentials provided when you created the PostgreSQL deployment.
 2. After signing in, open the navigation menu, and then click **Configuration**.
 3. For the PostgreSQL database connection, click **Connect**. Checkpoint table and TRANDATA fields appear if the connection is successful.
 4. Next to TRANDATA Information, click **Add TRANDATA** (plus icon).
 5. Enter a table name, schema name, or wildcard. For example, `src_ociggl1.*`.
 6. Click **Submit**.

 **Note:**

You only need to click Submit once. Use the search field to search for your table name and verify the tables were added.

Add an Extract

To add an Extract for PostgreSQL:

1. In the OCI GoldenGate deployment console, ensure that you're on the Administration Service Overview page, and then click **Add Extract** (plus icon).
2. For Extract Type, select one of the following:
 - Initial Load Extract
 - Change Data Capture Extract
3. Enter the Process Name, select the Credential Domain and Alias, and then enter a two-character name for the Trail name.
4. If you selected Initial Load Extract, click **Next** and skip to Step 7 - Extract Parameters. If you selected Change Data Capture Extract, you can click **Register only** to return to register the Extract and then return to the Overview page, or click **Next** to configure additional Extract Options.

 **Note:**

Register only registers the Extract without adding it. The registration creates the replication slot when you register the Extract or use the Register Only option.

5. On the Extract Options page, under Basic Information, complete the fields as needed:
 - a. For **Process Name**, enter a name for the Extract process, up to 8 characters.
 - b. For **Intent**, select the option that best describes the purpose of this Extract:
 - Now
 - Custom time
 - Position in log
 - End of log
 - c. For **Trail Name**, enter a two character name for the Trail file.
 - d. (Optional) For **Trail Subdirectory**, set a custom location for the generated Trail file.
 - e. (Optional) For **Trail Size**, set the max size for the generated trail file.
 - f. (Optional) Under Managed Options, you can configure the following:
 - Profile Name
 - Auto Start
 - Auto Restart

 **Note:**

Adding a profile and configuring Auto Start and Auto Restart options enables your deployment to restart automatically after a network disruption. See [Configure managed processes](#) for more information.

6. On the Extract Parameters page, you can edit the parameter file in the text area to list the table details to capture. For example:

```
table source.table1;
```

 **Note:**

GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")  
setenv (TZ="GMT+5")
```

[Learn more about SETENV.](#)

7. Click **Create and Run** to create and start the Extract. If you click **Create**, then you can manually start Extract later from the Administration Service Overview page.

You're returned to the Administration Service Overview page, where you can view the status of the Extract process. Select **Details** from the Extract **Action** menu to view process information, checkpoint, statistics, parameters, and reports.

Big Data Extracts

Learn to create Extracts for various Big Data technologies.

Add an Extract for Kafka

Adding an Extract for a Kafka source differs between Big Data deployment versions. Follow the instructions according to your Big Data deployment version.

To check your OCI GoldenGate version, view the Version information on the [deployment details](#) page.

You can capture messages from the following streaming sources:

- Apache Kafka
- OCI Streaming
- Confluent Kafka, with or without Schema Registry
- Azure Event Hubs
- Amazon MSK

OCI GoldenGate reads messages from a Kafka topic or topics, and then converts the data into logical change records written to GoldenGate Trail files. GoldenGate Replicat

processes can then use the generated Trail files to propagate data to support relational database management system (RDBMS) implementations.

For Big Data deployments version 21.8 or earlier

To add an Extract for Kafka in OCI GoldenGate for Big Data deployments version 21.8 or earlier:

1. Configure Consumer properties:

- a. Create a Kafka Consumer properties file with one of the following deserializers or converters. If the source is a topic in Confluent Kafka with Schema Registry, you can use the Avro converter. For other sources, use the JSON converter or deserializer as needed:

- Kafka Consumer properties for JSON deserializer:

```
key.deserializer=org.apache.kafka.common.serialization.ByteArrayDeseriali
zer
value.deserializer=org.apache.kafka.common.serialization.ByteArrayD
eserializer
```

- Kafka Consumer properties for JSON converter:

```
key.converter=org.apache.kafka.connect.json.JsonConverter
value.converter=org.apache.kafka.connect.json.JsonConverter
```

- Kafka Consumer properties for Avro converter:

```
key.converter=io.confluent.connect.avro.AvroConverter
value.converter=io.confluent.connect.avro.AvroConverter
```

- b. Save the properties file and note its location.
- c. Edit your Kafka connection and update the properties file in the **Show advanced options** section.

2. Create and run the Extract.

- a. On the Deployments page, select the Big Data deployment to view its details.
- b. On the deployment Details page, click **Launch console**.
- c. Log in to the deployment console using the credentials specified when you created the deployment.
- d. On the Administration Service Overview page, click **Add Extract** (plus icon).
- e. On the Add Extract page, select **Change Data Capture** for Extract type, and then click **Next**.
- f. On the Extract Options page, complete the fields as follows, and then click **Next**:
 - i. For **Process Name**, enter a name for the Extract.
 - ii. For **Alias**, select the Kafka connection assigned to the deployment.
 - iii. For **Begin**, select **Now**.
 - iv. For **Trail Name**, enter a 2-character name.
 - v. (Optional) **Enable Kafka Connect**, if the source is a Kafka Connect framework.

- vi. (Optional) Select a Converter. If you select **Avro**, select **Schema Registry**.
- vii. (Optional) Under Managed Options, you can configure the following:
 - Profile Name
 - Auto Start
 - Auto Restart

 **Note:**

Adding a profile and configuring Auto Start and Auto Restart options enables your deployment to restart automatically after a network disruption. See Configure managed processes for more information.

- g. On the Parameter file page, leave the table mapping as is (`TABLE TESTSCHEMA.*;`), to listen to all topics in the given bootstrap server. To capture from a designated topic, change the mapping to `TABLE TESTSCHEM.<topic-name>;` where `<topic-name>` is the name of the topic to capture from.
- h. Click **Create and Run**.

You return to the Administration Service Overview page where you can observe the Extract status. You can then select the Extract to view its details, statistics, and reports.

For Big Data deployments 21.9 or later

OCI GoldenGate Big Data deployment versions 21.9 and later supports multiple source technologies within the same deployment, so you can configure multiple extracts from different supported source technologies within the same deployment. OCI GoldenGate uses credential store entries to identify the source technology. For example, OCI GoldenGate uses the credential's *User ID* to determine whether to start a Kafka (`kafka://`) or Mongo (`Mongo://`) Extract.

Before you create an Extract, create a credential in the Big Data deployment console:

1. On the Deployments page, select the Big Data deployment to view its details.
2. Log in to the Big Data deployment console using the credentials specified when you created the deployment. You're brought to the Administration Server Overview page.
3. Add a credential:
 - a. Open the navigation menu, and then click **Configuration**.
 - b. On the Credentials page, click **Add Credential** (plus icon).
 - c. Enter the following details in the fields provided, and then click **Submit**:
 - For **Credential Domain**, enter `OracleGoldenGate`.
 - For **Credential Alias**, enter `kafka`.
 - For **User ID**: enter `kafka://`
 - For **Password**, enter a password.

- For **Verify Password**, enter the password again.

To add an Extract for Kafka in OCI GoldenGate for Big Data deployments version 21.9 or later:

1. In the navigation menu, click **Overview** to return to the Administration Service Overview page.
2. Click **Add Extract** (plus icon).
3. On the Add Extract page, select the following, and then click **Next**:
 - Source: **Kafka**
 - Extract type: **Change Data Capture Extract**
4. On the Extract Options page, complete the fields as follows, and then click **Next**:
 - For Process Name, enter a name, up to 8 characters.
 - (Optional) Enter a description.
 - For Connection Alias, select the previously assigned connection alias.
 - For Begin, select **Now**.
 - For Trail Name, enter a 2-character name.
 - (Optional) **Enable Kafka Connect**, if the source is a Kafka Connect framework.
 - (Optional) If Kafka Connect is selected, select one of the available Converters:
 - JSON
 - AVRO
 - (Optional) If AVRO is selected, select **Schema Registry**.
 - (Optional) Under Managed Options, you can configure the following:
 - Profile Name
 - Auto Start
 - Auto Restart

 **Note:**

Adding a profile and configuring Auto Start and Auto Restart options enables your deployment to restart automatically after a network disruption. See Configure managed processes for more information.

5. On the Parameter file page, update the following:
 - a. Update `SOURCEDB USERIDALIAS` to `SOURCEDB USERIDALIAS kafka DOMAIN OracleGoldenGate`
 - b. Leave the table mapping as is (`TABLE source.*;`), to listen to all topics in the given bootstrap server. To capture from a designated topic, change the mapping to `TABLE source.<topic-name>;` where `<topic-name>` is the name of the topic to capture from.

 **Note:**

The Extract mapping format is `Table SourceSchema.Table`. In Kafka, there are no source schemas/tables. OCI GoldenGate writes the first part ("source" in above step) as the schema name and the second part as table name to trail file. This way, you can replicate the captured kafka message into other GoldenGate supported targets.

6. Click Create and Run.

You return to the Administration Service Overview page where you can observe the Extract status. You can then select the Extract to view its details, statistics, and reports.

Add an Extract for MongoDB

Extract is a process that runs against the source data source connection and extracts, or captures, data. Learn to add an Extract for MongoDB.

Add an Extract for MongoDB

Learn to add an Extract process for a MongoDB source in OCI GoldenGate

To add an Extract for MongoDB:

1. In the OCI GoldenGate deployment console, ensure that you're on the Administration Service Overview page, and then click **Add Extract** (plus icon).
2. On the Add Extract page, for Extract type select one of the following, and then click **Next**:
 - Initial Load
 - Change Data Capture
3. On the Extract Options page, complete the fields as follows, and then click **Next**:
 - a. For **Process Name**, enter a name for the Extract.
 - b. For **Connection Alias**, select the connection alias from the dropdown.
 - c. For **Source**, select **File**.
 - d. For **File Name**, enter three characters at minimum for the filename.
 - e. (Optional) Under Managed Options, you can configure the following:
 - Profile Name
 - Auto Start
 - Auto Restart

 **Note:**

Adding a profile and configuring Auto Start and Auto Restart options enables your deployment to restart automatically after a network disruption. See Configure managed processes for more information.

4. On the Parameter File page, update the source mapping with `TABLE source.*;`

 **Note:**

This is the source database/collection mapping. `TABLE *.*`; results in the process extracting from all databases/collections.

5. Click Create and Run.

You return to the Administration Service Overview page. Click the Extract name to view details and reports of the Extract.

Known issues

Replicats fail when using Trail file from MongoDB Extract with BINARY_JSON_FORMAT

When a Replicat uses a Trail file generated from a MongoDB Extract with BINARY_JSON_FORMAT in the Extract parameter file, the Replicat fails with the following error:

```
ERROR 2023-08-04 17:13:13.000421 [main] - Unable to decode column 0 : Input length = 1
```

```
java.nio.charset.MalformedInputException: Input length = 1 at  
java.nio.charset.CoderResult.throwException(CoderResult.java:281)
```

```
~[?:1.8.0_311]at
```

```
java.nio.charset.CharsetDecoder.decode(CharsetDecoder.java:816)
```

```
~[?:1.8.0_311] at
```

```
oracle.goldengate.datasource.UserExitDataSource.createColumnValue(UserExitData  
aSource.java:1106)
```

```
[ggdbutil-21.9.0.0.3.001.jar:21.9.0.0.3.001] Exception in thread "main"  
oracle.goldengate.util.GGException: Unable to decode column 0 : Input  
length = 1 at
```

```
oracle.goldengate.datasource.UserExitDataSource.createColumnValue(UserExitDat  
aSource.java:1203)
```

Workaround: When BINARY_JSON_FORMAT is removed from the Extract parameters, the Replicat runs successfully and documents are represented in Extended JSON format.

Add a Distribution Path

A Distribution path sends the transaction of data from an Extract to a Replicat.

When to use a Distribution Path

Use a Distribution Path when you need to replicate data in a distributed deployment environment. A Distribution Path sends the transaction of data from the Extract to the Replicat. Creating and running a Distribution Path automatically creates a Receiver Path in the target deployment's Receiver service. The Receiver Path receives the transaction of data from the source deployment's Distribution service.

Before you begin

Ensure that you create GoldenGate connections for each deployment you want to connect to, and then assign them to the deployment from which the path originates. For example, the deployment where you create the Distribution Path or target-initiated Receiver Path.

Otherwise, in IAM-enabled deployments, you'll encounter the error:

```
The network connection could not be established: 'OGG-08654' -  
'Invalid or missing OAuth  
resource - audiencescope in Client application'.
```

Create and run a Distribution Path


To add a Distribution Path:

1. In the OCI GoldenGate deployment console, click **Distribution Service**.
2. On the Distribution Service Overview page, click **Add Path** (plus icon).
3. On the Add Path page, complete the fields as follows:
 - a. For **Path Name**, enter a name for the path.
 - b. (Optional) For **Description**, enter a brief description of this path's purpose.
 - c. (Optional) Enable **Reverse proxy enabled** to connect to the target using a reverse proxy.
 - d. For **Source**, select the Extract name from the dropdown.
 - e. For **Trail Name**, select the Extract trail from the dropdown.
 - f. For **Generated Source URI**, the URI is automatically generated based on the Extract information provided. Click Edit (pencil icon) to modify the URI, if needed.
 - g. For **Target Authentication Method**, select the authentication method for the target URI:
 - **OAuth**: Select this option if the source and target deployments are IAM enabled. This option uses the client credentials for authentication from the Distribution Service to Receiver Service.

 **Note:**

In IAM enabled tenancies, select **OAuth** when connecting to another IAM-enabled deployment.

- **UserID Alias**: This option uses a UserID Alias that you can create on the target Oracle GoldenGate to establish a connection.

 **Note:**

Create a credential using the IAM user when connecting to an IAM-enabled deployment from a GoldenGate credential store deployment. Ensure that the user exists in the IAM Identity Stripe.

- h. For **Target**, select a data transfer protocol from the dropdown:
- **wss**: Web socket secure is the default option. If selected, you must complete the following fields:
 - Target Host
 - Port Number
 - Trail Name
 - Deployment Name (reverse proxy enabled)
 - URI Path (reverse proxy enabled)
 - Domain
 - Alias
 - **ogg**: If select, you must complete the following fields:
 - Target Host
 - Port Number
 - Trail Name
 - Deployment Name (reverse proxy enabled)
 - URI Path (reverse proxy enabled)
 - **ws**: If selected, you must complete the following fields:
 - Target Host
 - Port Number
 - Trail Name
 - Deployment Name (reverse proxy enabled)
 - URI Path (reverse proxy enabled)
 - Domain
 - Alias
- i. For **Generated Target URI**, the URI is automatically populated from the information provided. Click Edit (pencil icon) to modify the URI, if needed.
- j. For Target Encryption Algorithm, select an encryption algorithm for the target trail:
- None
 - AES128
 - AES192
 - AES256
- k. For **Enable Network Compression**, select this option to set the Compression Threshold.

- l. For **Sequence Length**, enter the length of the trail sequence number.
- m. For **Trail Size (MB)**, enter the maximum size for a file in a trail.
- n. For **Configure Trail Format**, enable this option if you want to configure the trail file format, and then complete the additional fields as needed.
4. Under the Encryption Profile section, complete the following fields as needed:
 - a. Profile Name
 - b. Encryption Profile Type
 - c. Masterkey Name
 - d. For **Begin**, select where to log data:
 - Now
 - Custom Time
 - Position in Log (default)
 - e. For **Source Sequence Number**, select the sequence number of the trail file source deployment Extract.
 - f. For **Source RBA Offset**, enter the Relative Byte Address (RBA) in the trail file where you want the process to start.
 - g. For **Critical**, set this option to True if the distribution path is critical to the deployment. The default is False.
 - h. For **Auto Restart**, set this option to True if you want the distribution path to restart automatically if it's terminated.
 - i. For **Auto Restart Options**, indicate the number of retries to restart the path process and the delay duration interval between retries.
5. Under Rule-set Configuration, complete the following fields as needed:
 - For **Enable Filtering**, if selected, click **Add Rule**, and then complete the additional fields.
6. Under More Options, complete the following fields as needed:
 - a. EOF Delay: end of file delay before searching for source data
 - b. Checkpoint Frequency: frequency in seconds for routine checkpoints
 - c. App Options
 - TCP Flush Bytes: Flush size
 - TCP Flush Seconds: Flush interval
 - d. TCP Options
 - DSCP: network differentiated services
 - TOS term of service
 - TCP_NODELAY: disables use of Nagle's algorithm if enabled
 - Quick ACK: sends acknowledgement if enabled
 - TCP_CORK: enables use of Nagle's algorithm
 - System Send buffer Size
 - System Receive Buffer Size

- Keep Alive: timeout for keep alive

7. Click **Create and Run**.

You return to the Overview page where you can view the status of the Path process.

Learn more

Interested in learning more about the Distribution Path process? Refer to the following resources:

- [Send data from Oracle GoldenGate to OCI GoldenGate](#)
- [Replicate data from MySQL to Autonomous Data Warehouse](#)
- [Replicate Autonomous Transaction Processing to OCI Object Storage](#)

Known issues

To create Distribution Paths to send data to or pull data from Oracle Cloud Infrastructure GoldenGate, ensure that you add the root certificate to Certificate Management or your client wallet

To send data to or pull data from OCI GoldenGate, you must create a Distribution Server Path or a target initiated path on the Receiver Server in your on-premises or Marketplace Oracle GoldenGate, respectively. You must also add the OCI GoldenGate root certificate or self-signed certificate to your Oracle GoldenGate Certificate Management (Oracle GoldenGate 21c or higher) or client wallet (Oracle GoldenGate 19c). This creates a trusted connection between your Oracle GoldenGate and OCI GoldenGate deployments. Only WebSocket Secure (WSS) protocol is supported for Distribution and Receiver Server Paths between Oracle GoldenGate and OCI GoldenGate.

A change in the OCI GoldenGate root certificate will cause the Distribution Server Path or a target initiated path on the Receiver Server in your on-premises or Marketplace Oracle GoldenGate to fail and produce the following error:

```
ERROR   OGG-10390  Oracle GoldenGate Receiver Service:  Generic error -1
noticed for endpoint
        wss://<deployment URL>:443/services/v2/sources?trail=<trail name>.
Error description - SSL
        connection unexpectedly closed.
```

Workaround: To fix this issue, update the certificate in the client wallet or Service Manager's Certificate Management screen to use the OCI GoldenGate Deployment Console root certificate.

Learn more:

- For Oracle GoldenGate 19c users, see [Creating a Distribution Server Path User Certificate](#).
- For users of Oracle GoldenGate 21c or higher, see [Create a Trusted Connection Between Oracle GoldenGate and OCI GoldenGate](#).

Only Digest Authentication is currently supported

Oracle Cloud Infrastructure GoldenGate doesn't currently support certificate-based authentication when you use Oracle Cloud Infrastructure GoldenGate as the Distribution Path target.

Workaround: None.

Add a Receiver Path

A Receiver path handles incoming trail files.

When to use target-initiated Receiver Paths

Use target-initiated Receiver Paths when the network security policies prevent the source deployment's Distribution service from opening network connections to the target environment's Receiver service. If a source deployment's Distribution server is unable to initiate connections to the target deployment's Receiver service, but the Receiver service can initiate a connection to the source deployment's Distribution service, then you can create and run a target-initiated Receiver path to pull Trail files from the source Distribution service.

Before you begin

Ensure that you create GoldenGate connections for each deployment you want to connect to, and then assign them to the deployment from which the path originates. For example, the deployment where you create the Distribution Path or target-initiated Receiver Path.

Otherwise, in IAM-enabled deployments, you'll encounter the error:

```
The network connection could not be established: 'OGG-08654' -  
'Invalid or missing OAuth  
resource - audiencescope in Client application'.
```

Create and run a Receiver Path

To add a Receiver path:

1. In the OCI GoldenGate deployment console, click **Receiver Service**.
2. On the Overview page, click **Add Path** (plus icon).
3. On the Add Path page, complete the fields as follows:
 - a. For **Path Name**, enter a name for the path.
 - b. (Optional) For **Description**, enter a brief description of this path's purpose.
 - c. (Optional) Enable **Reverse proxy enabled** to connect to the target using a reverse proxy.
 - d. For **Source Authentication Method**, select one of the following:

- **OAuth:** Select this option if the source and target deployments are IAM with Identity Domains enabled. This option uses the client credentials for authentication from the Distribution Service to Receiver Service.

 **Note:**

In IAM enabled tenancies, select **OAuth** when connecting to another IAM-enabled deployment.

- **UserID Alias:** This option uses a UserID Alias that you can create on the target Oracle GoldenGate to establish a connection.

 **Note:**

Create a credential using the IAM user when connecting to an IAM-enabled deployment from a GoldenGate credential store deployment. Ensure that the user exists in the IAM Identity Stripe.

- e. For **Source**, select a data transfer protocol from the dropdown:
 - **ws:** If selected, Source Authentication Method is not required.
 - **wss:** Web socket secure is the default option.
- f. For **Source Host**, enter the host name of the source deployment.
- g. For **Port Number**, enter the source deployment port number.
- h. For **Trail Name**, enter the name of the source deployment's Extract trail file.
- i. **Deployment Name** is required if Reverse Proxy is enabled. Enter the source deployment name.
- j. **URI Path** is required if Reverse Proxy is enabled. Enter the source deployment URI path.
- k. For **Domain**, enter the source credential domain.
- l. For **Alias**, enter the source credential alias.
- m. **Generated Source URI** autopopulates based on the values entered. Click **Edit** (pencil icon) to modify the URI, if needed.
- n. For **Target Trail Name**, enter the name of the Trail file as it's received by the target.
- o. For Target Encryption Algorithm, select an encryption algorithm for the target trail:
 - None
 - AES128
 - AES192
 - AES256
- p. For **Enable Network Compression**, select this option to set the Compression Threshold.
- q. For **Sequence Length**, enter the length of the trail sequence number.
- r. For **Trail Size (MB)**, enter the maximum size for a file in a trail.

- s. For **Configure Trail Format**, enable this option if you want to configure the trail file format, and then complete the additional fields as needed.
4. Under the Encryption Profile section, complete the following fields as needed:
 - a. Profile Name
 - b. Encryption Profile Type
 - c. Masterkey Name
 - d. For **Begin**, select where to log data:
 - Now
 - Custom Time
 - Position in Log (default)
 - e. For **Source Sequence Number**, select the sequence number of the trail file source deployment Extract.
 - f. For **Source RBA Offset**, enter the Relative Byte Address (RBA) in the trail file where you want the process to start.
 - g. For **Critical**, set this option to True if the distribution path is critical to the deployment. The default is False.
 - h. For **Auto Restart**, set this option to True if you want the distribution path to restart automatically if it's terminated.
 - i. For **Auto Restart Options**, indicate the number of retries to restart the path process and the delay duration interval between retries.
5. Under Rule-set Configuration, complete the following fields as needed:
 - For **Enable Filtering**, if selected, click **Add Rule**, and then complete the additional fields.
6. Under More Options, complete the following fields as needed:
 - a. EOF Delay: end of file delay before searching for source data
 - b. Checkpoint Frequency: frequency in seconds for routine checkpoints
 - c. App Options
 - TCP Flush Bytes: Flush size
 - TCP Flush Seconds: Flush interval
 - d. TCP Options
 - DSCP: network differentiated services
 - TOS term of service
 - TCP_NODELAY: disables use of Nagle's algorithm if enabled
 - Quick ACK: sends acknowledgement if enabled
 - TCP_CORK: enables use of Nagle's algorithm
 - System Send buffer Size
 - System Receive Buffer Size
 - Keep Alive: timeout for keep alive
7. Click **Create and Run**.

You return to the Overview page, where you can view the status of the path process.

Learn more

Interested in learning how to create and run a target-initiated Receiver Path? Refer to the [Send data from OCI GoldenGate to Oracle GoldenGate quickstart](#) or try the [LiveLabs version in a sandbox environment](#).

Add Replicats

Learn to configure Replicats for various types of target connections.

Relational Database Management System (RDBMS) Replicats:

- [Add a Repicat for Oracle Database](#)
- [Add a Repicat for Microsoft SQL Server](#)
- [Add a Repicat for MySQL](#)
- [Add a Repicat for PostgreSQL](#)

Big Data Replicats:

- [Add a Repicat for Autonomous Database](#)
- [Add a Repicat for Autonomous JSON Database](#)
- [Add a Repicat for OCI Object Storage](#)
- [Add a Repicat for OCI Streaming](#)
- [Add a Repicat for Amazon Kinesis](#)
- [Add a Repicat for Amazon Redshift](#)
- [Add a Repicat for Amazon S3](#)
- [Add a Repicat for Azure Cosmos DB for MongoDB](#)
- [Add a Repicat for Azure Data Lake Storage](#)
- [Add a Repicat for Azure Event Hubs](#)
- [Add a Repicat for Azure Synapse Analytics](#)
- [Add a Repicat for Kafka](#)
- [Add a Repicat for Confluent Kafka](#)
- [Add a Repicat for Elasticsearch Server](#)
- [Add a Repicat for Google BigQuery](#)
- [Add a Repicat for Google Cloud Storage](#)
- [Add a Repicat for MongoDB](#)
- [Add a Repicat for Redis](#)
- [Add a Repicat for Snowflake](#)

RDBMS Replicats

Learn to add Replicats for various target relational database management systems (RDBMS), such as Oracle Database, Microsoft SQL Server, MySQL, and PostgreSQL.

Add a Repicat for Oracle Database

A Repicat is a process that delivers data to a target database or technology. It reads the trail file on the target, reconstructs the DML or DDL operations, and then applies them to the target. Learn to add a Repicat for Oracle Database, OCI Autonomous Databases, Oracle Exadata, and Amazon RDS for Oracle technologies.

Create a Checkpoint table

Before you add a Repicat, create a checkpoint table to ensure the Repicat can restart without reapplying transactions should a disruption occur.

To create a checkpoint table:

1. In the OCI GoldenGate deployment console navigation menu, click **Configuration**.
2. In the **Administration Service** tab, go to the **Database** tab.
3. Click the database connect icon for the target database.
4. Click **Add Checkpoint** (plus icon).
5. For Checkpoint Table, enter the target schema name, and then click **Submit**.
6. Reconnect to the target database.

Add a Repicat

To add a Repicat for Oracle Database:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Repicat** (plus icon).
2. On the Add Repicat page, select a Repicat type, and then click **Next**.

The Repicat types are:

- Parallel Repicat
- Integrated Repicat
- Coordinated Repicat
- Classic Repicat

For more information about the different Repicat types, see [Types of Repicat](#).

3. On the Repicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Repicat process.
 - b. (Optional) For **Description**, enter a brief description to distinguish this Repicat from others.
 - c. For **Intent**, select the Repicat's purpose:

- High Availability
 - Disaster Recovery
 - Unidirectional (default)
 - N-Way
- d. For **Credential Domain** and **Credential Alias**, select the target database domain and alias, or create a new credential.
 - e. For **Source**, select the source of data to process, either **Trail** or **File**.
 - f. If the source is Trail, then for **Trail Name**, enter the name of the Extract trail.
 - g. If the source is Trail and you expect the trail file to be in a location other than the default, then enter the location for **Trail Subdirectory**.
 - h. For **Begin**, select the starting point for data processing:
 - Position in Log
 - Now
 - Custom Time
 - i. For **Transaction Log Sequence Number**, leave the default value or enter a transaction log sequence number.
 - j. For **Transaction Log RBA Offset**, leave the default value or enter an offset value.
 - k. For **Checkpoint Table**, leave the default selection or select the checkpoint table you created for the target deployment.
4. Under **Encryption Profile**, select a profile name. If an encryption profile wasn't created, the Local Wallet is selected by default.
 - a. Encryption Profile Type
 - b. Masterkey Name
 5. Under **Managed Options**, select **Critical to deployment health**, and then completed the other fields as needed.
 6. Click **Next**.
 7. On the Parameter File page, you can specify parameters to further configure your Replicat.

```
table source.table1;
```

 **Note:**

GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")  
setenv (TZ="GMT+5")
```

Learn more about SETENV.

8. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Microsoft SQL Server

Learn to add and configure a Replicat process for Microsoft SQL Server, Amazon RDS for SQL Server, Azure SQL Database, and Azure SQL Managed Instance targets.

To add a Replicat for Microsoft SQL Server:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, under **Replicat type**, select from the following options, and then click **Next**:
 - Classic Replicat
 - Coordinated Replicat
 - Parallel Replicat

 **Note:**

Learn about Replicat types.

3. On the Replicat Options page, under Basic Information, complete the following fields as needed:
 - a. For **Process Name**, enter a name for the Replicat
 - b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.
 - c. For **Intent**, select the Replicat's purpose:
 - High Availability
 - Disaster Recovery
 - Unidirectional (default)
 - N-Way
 - d. For **Credential Domain** and **Credential Alias**, select the target database domain and alias, or create a new credential.
 - e. For **Source**, select the source of data to process, either **Trail** or **File**.
 - f. If the source is Trail, then for **Trail Name**, enter the name of the Extract trail.
 - g. If the source is Trail and you expect the trail file to be in a location other than the default, then enter the location for **Trail Subdirectory**.
 - h. For **Begin**, select the starting point for data processing:
 - Position in Log
 - Now
 - Custom Time

- i. For **Trail Sequence Number**, leave the default value or enter a trail sequence number.
 - j. For **Trail RBA Offset**, leave the default value or enter a trail offset.
 - k. For **Checkpoint Table**, leave the default selection or select the checkpoint table you created for the target deployment.
4. Click **Next**.
 5. For **Parameter File**, you can add parameters to further configure the Replicat process.

 **Note:**

GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

[Learn more about SETENV.](#)

6. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You're returned to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for MySQL

Learn to add and configure a Replicat process for MySQL Database Server, OCI MySQL Heatwave, Amazon Aurora MySQL, Amazon RDS for MySQL, Amazon RDS for MariaDB, Azure Database for MySQL, and Google Cloud SQL for MySQL target technologies.

To add a Replicat for MySQL:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, under **Replicat type**, select from the following options, and then click **Next**:
 - Classic Replicat
 - Coordinated Replicat
 - Parallel Replicat

 **Note:**

See [Decide which apply method to use](#) to learn more about Replicat types.

3. On the Replicat Options page, under Basic Information, complete the following fields as needed:
 - a. For **Process Name**, enter a name for the Replicat

- b. (Optional) For **Description**, enter a brief description to distinguish this Repicat from others.
 - c. For **Intent**, select the Repicat's purpose:
 - High Availability
 - Disaster Recovery
 - Unidirectional (default)
 - N-Way
 - d. For **Credential Domain** and **Credential Alias**, select the target database domain and alias, or create a new credential.
 - e. For **Source**, select the source of data to process, either **Trail** or **File**.
 - f. If the source is Trail, then for **Trail Name**, enter the name of the Extract trail.
 - g. If the source is Trail and you expect the trail file to be in a location other than the default, then enter the location for **Trail Subdirectory**.
 - h. For **Begin**, select the starting point for data processing:
 - Position in Log
 - Now
 - Custom Time
 - i. For **Trail Sequence Number**, leave the default value or enter a trail sequence number.
 - j. For **Trail RBA Offset**, leave the default value or enter a trail offset.
 - k. For **Checkpoint Table**, leave the default selection or select the checkpoint table you created for the target deployment.
4. Click **Next**.
 5. For **Parameter File**, you can add parameters to further configure the Repicat process.

 **Note:**

GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

6. Click **Create and Run**. If you click **Create**, then you can manually start the Repicat later from the Administration Service Overview page.

You're returned to the Administration Service Overview page where you can view the creation of the Repicat process and access Repicat Actions.

Add a Replicat for PostgreSQL

To add a Replicat process for a PostgreSQL target:

1. In the deployment console, on the Administration service Overview page, click **Add Replicat** (plus icon)
2. In the Add Replicat wizard, on the Replicat Type page, select one of the following, and then click **Next**:
 - Classic Replicat
 - Coordinated Replicat
 - Parallel Replicat

Learn more about different Replicat types.

3. On the Replicat Options page, complete the following fields:
 - a. For **Process Name**, enter name, no more than 5 characters long.
 - b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.
 - c. For **Intent**, select the Replicat's purpose:
 - Disaster Recovery
 - Unidirectional (default)
 - N-Way
 - d. For **Credential Domain** and **Credential Alias**, select the target database domain and alias, or create a new one.
 - e. For **Source**, select the source of data to process, either **Trail** or **File**.
 - f. If the source is Trail, then for **Trail Name**, enter the name of the Extract trail. If the source is File, then enter the **File Name**.
 - g. (Optional) If the source is Trail and you expect the trail file to be in a location other than the default, then enter the location for **Trail Subdirectory**.
 - h. For **Begin**, select the starting point for data processing:
 - Position in Log (default)
 - Now
 - Custom Time
 - i. For **Trail Sequence Number**, leave the default value or enter a transaction log sequence number.
 - j. For **Trail RBA Offset**, leave the default value or enter an offset value.
 - k. For **Checkpoint Table**, select the checkpoint table you created for the target deployment, or enter a two part checkpoint table name.
4. Under Encryption Profile, select a profile name. If an encryption profile wasn't created, the Local Wallet is selected by default.
5. Click **Next**.
6. On the Parameter File page, add parameters to further configure the Replicat process.

 **Note:**

GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")  
setenv (TZ="GMT+5")
```

[Learn more about SETENV.](#)

7. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Big Data Replicats

Learn to add Replicats for target Big Data connections.

Add a Replicat for Autonomous Database

To add a Replicat for Autonomous Database:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The Replicat types are:

- Parallel Replicat
- Integrated Replicat
- Coordinated Replicat
- Classic Replicat

For more information about the different Replicat types, see [Types of Replicat](#).

3. On the Replicat Options page, under Basic Information, complete the following fields as needed:
 - a. For **Process Name**, enter a name for the Replicat
 - b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.
 - c. For **Intent**, select the Replicat's purpose:
 - High Availability
 - Disaster Recovery
 - Unidirectional (default)
 - N-Way
 - d. For **Credential Domain** and **Credential Alias**, select the target database domain and alias, or create a new credential.

- e. For **Source**, select the source of data to process, either **Trail** or **File**.
 - f. If the source is Trail, then for **Trail Name**, enter the name of the Extract trail.
 - g. If the source is Trail and you expect the trail file to be in a location other than the default, then enter the location for **Trail Subdirectory**.
 - h. For **Begin**, select the starting point for data processing:
 - Position in Log
 - Now
 - Custom Time
 - i. For **Transaction Log Sequence Number**, leave the default value or enter a transaction log sequence number.
 - j. For **Transaction Log RBA Offset**, leave the default value or enter an offset value.
 - k. For **Checkpoint Table**, leave the default selection or select the checkpoint table you created for the target deployment.
4. Under **Encryption Profile**, select a profile name. If an encryption profile wasn't created, the Local Wallet is selected by default.
 - a. Encryption Profile Type
 - b. Masterkey Name
 5. Under **Managed Options**, select **Critical to deployment health**, and then completed the other fields as needed.
 6. Click **Next**.
 7. On the Parameter Files page, you can specify parameters to further configure your Replicat.

```
table source.table1;
```

 **Note:**

- If using Coordinated Replicat for an Autonomous Data Warehouse target, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/
Deployment/etc/conf/ogg/<replicat
name>.properties
```

If omitted, you will encounter the following error:

```
OGG-01091 Oracle GoldenGate Delivery, RSNOW.prm: Unable
to open file
"/u02/Deployment/etc/conf/ogg/<replicat
name>001.properties" (error 2, No such file or
directory).
```

- GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

8. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You're returned to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Oracle Autonomous JSON Database

Learn to add a Replicat process for an Oracle Autonomous JSON Database target.

To add a Replicat for an Oracle Autonomous JSON Database target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The types of Replicats are:

- Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.

- c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **ORACLE_AUTONOMOUS_JSON_DATABASE** from the dropdown.
 - e. For **Available aliases**, select the Oracle Autonomous JSON Database connection.
4. On the Parameter File page, add and configure Replicat parameters as needed, such as `MAP *.*`, `TARGET *.*`; and then click **Next**.

 **Note:**

- If using Coordinated Replicat, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/
Deployment/etc/conf/ogg/<replicat
name>.properties
```

Using coordinated replicat will result in multiple files created.

If omitted, you will encounter the following error:

```
OGG-01091 Oracle GoldenGate Delivery, RSNOW.prm: Unable to
open file
"/u02/Deployment/etc/conf/ogg/<replicat
name>001.properties" (error 2, No such file or
directory).
```

- GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for an Object Storage target

Learn to add and configure a Replicat process for an OCI Object Storage target.

Before you begin

Before you add the Replicat, ensure that you have the following:

- Compartment OCID

 **Tip:**

You can find the compartment OCID on the Compartment details page.

1. Open the Oracle Cloud console navigation menu, select **Identity**, and then **Compartmentments**.
2. Select your compartment from the list to access the Compartment details page.
3. Copy the compartment OCID from the Compartment Information section.

For more information, see [Find the OCID of a Compartment](#).

- Oracle Object Storage bucket name

 **Tip:**

Oracle recommends that you create your own Object Storage bucket.

1. Open the Oracle Cloud console navigation menu, select **Storage**, and then **Buckets**.
2. Click **Create Bucket**.
3. In the Create Bucket panel, enter a name, and then click **Create**.

For more information, see [Using the Console to create a bucket](#).

Add a Replicat

To add a Replicat for an Object Storage target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The types of Replicats are:

- Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For Process Name, enter a name for the Replicat process.
 - b. (Optional) For Description, enter a short description to distinguish this process from others.
 - c. For Trail Name, enter two-character trail name.
 - d. For Target, select **Oracle Object Storage** from the dropdown.
 - e. For Available aliases for OCI, select your alias from the dropdown.
 4. On the Parameter File page, add and configure Replicat parameters as needed, and then click **Next**.

 **Note:**

- If using Coordinated Replicat, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/
Deployment/etc/conf/ogg/<replicat
name>.properties
```

Using coordinated replicat will result in multiple files created.

If omitted, you will encounter the following error:

```
OGG-01091 Oracle GoldenGate Delivery, RSNOW.prm: Unable to
open file
"/u02/Deployment/etc/conf/ogg/<replicat
name>001.properties" (error 2, No such file or
directory).
```

- GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. On the Properties File page, configure the File Handler and OCI Event Handler properties as needed, and then click **Next**. Some properties to consider modifying include:
 - a. `gg.handler.name.format`: Select how to format the output data. Available options include:
 - `delimitedtext`
 - `json`
 - `json_row`
 - `xml`
 - `avro_row`
 - `avro_op`
 - `avro_row_ocf`
 - `avro_op_ocf`
 - b. `gg.handler.name.fileNameMappingTemplate`: Generates file names dynamically using Template Keywords.
 - c. `gg.handler.name.inactivityRollInterval`: GoldenGate creates a file and keeps it open for writing. This property closes the file after the designated period of inactivity (no incoming transactions) and then loaded into OCI Object Storage.
 - d. `gg.eventhandler.name.compartmentID`: Enter the compartment OCID.

- e. `gg.eventhandler.name.bucketMappingTemplate`: Enter the Object Storage bucket name.

Learn more about File Writer Handler and OCI Event Handler properties.

 **Note:**

You can also add pluggable formatters as needed. For more information, see [Using the Pluggable Formatters](#).

6. Click **Add and Run**.

You're returned to the Administration Service Overview page where you can monitor the status of the Replicat process. Click the process name to view its details and access reports.

Add a Replicat for an OCI Streaming target

Learn to add and configure a Replicat process for an OCI Streaming target.

Before you begin

Before you add the Replicat, ensure that you have the following:

- An Auth Token

 **Tip:**

Create the Auth Token on your user details page.

1. In the Oracle Cloud console global navigation bar, click **Profile**, and then select your username.
2. On your User details page, under **Resources**, click **Auth Tokens**.
3. Click **Generate Token**.
4. In the Generate Token dialog, enter a description, and then click **Generate Token**.
5. Copy and paste your token to a text editor.

- [A stream created in OCI Streaming](#).

Add a Replicat

To add a Replicat for an OCI Streaming target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The types of Replicats are:

- Classic Replicat

- Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a short description to distinguish this process from others.
 - c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **OCI Streaming** from the dropdown.
 - e. For **Available aliases for OCI**, select the OCI Streaming connection.
 4. On the Parameter File page, add and configure Replicat parameters as needed, and then click **Next**.

 **Note:**

GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

[Learn more about SETENV.](#)

5. On the Properties File page, for `gg.handler.kafkahandler.topicMappingTemplate=`, enter the stream name.
6. Click **Add and Run**.

You're returned to the Administration Service Overview page where you can monitor the status of the Replicat process. Click the process name to view its details and access reports.

Add a Replicat for Amazon Kinesis

Learn to create a Replicat process for an Amazon Kinesis target in OCI GoldenGate.

Before you begin

Before adding and running a Replicat for Amazon Kinesis, ensure that you have the following:

- An Amazon Kinesis connection created and assigned to your target Big Data deployment
- Your Amazon Web Services (AWS) region

Add a Replicat for Amazon Kinesis

Learn to add a replicat process for Amazon Kinesis.

To add a Replicat for an Amazon Kinesis target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The types of Replicats are:

- Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.
 - c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **Amazon Kinesis** from the dropdown.
 - e. For **Available Aliases for Amazon Kinesis**, select your alias from the dropdown.
 4. On the Parameter File page, add and configure Replicat parameters as needed, and then click **Next**.

 **Note:**

- If using Coordinated Replicat, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/  
Deployment/etc/conf/ogg/<replicat  
name>.properties
```

Using coordinated replicat will result in multiple files created.

If omitted, you will encounter the following error:

```
OGG-01091 Oracle GoldenGate Delivery, RSNOW.prm: Unable  
to open file  
"/u02/Deployment/etc/conf/ogg/<replicat  
name>001.properties" (error 2, No such file or  
directory).
```

- GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")  
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. On the Properties File page, configure the following Amazon Kinesis properties:
 - `gg.eventhandler.kinesis.region::` provide the AWS region for the target Amazon Kinesis stream.

- `gg.handler.kinesis.streamMappingTemplate=${tableName}`, **replace** `${tableName}` with the name of an existing data stream, leave the default value as is, or use Template Keywords to assign the stream name dynamically.
 - (Optional) `gg.handler.kinesis.partitionMappingTemplate=${primaryKeys}`, leave the default value as is, or replace `${primaryKeys}` with the template name to resolve the message key.
6. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Amazon Redshift

Learn to create a Replicat process for an Amazon Redshift target in OCI GoldenGate.

Before you begin

Before adding and running a Replicat for Amazon Redshift, ensure that you have the following:

- Amazon Redshift and Amazon S3 connections created and assigned to your target deployment.
- An Amazon Redshift cluster.
- Target tables created in Amazon Redshift.
- An Amazon S3 bucket configured in the same region with your Amazon Redshift cluster.
- An [IAM role](#) created to allow Amazon Redshift to access Amazon S3 buckets.

Add a Replicat for Amazon Redshift

Learn to add a replicat process for Amazon Redshift.

To add a Replicat for an Amazon Redshift target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The types of Replicats are:

- Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.
 - c. If the source is Trail, then for **Trail Name**, enter the name of the Extract trail.
 - d. For **Target**, select Amazon Redshift from the dropdown.
 - e. For **Available aliases** for Amazon Redshift, select your alias from the dropdown.
 - f. For **Available staging locations**, select Amazon S3 from the dropdown.

- g. For via staging alias, select Amazon S3 connection from the dropdown.
4. On the Parameter File page, you can specify parameters to further configure your Replicat.

 **Note:**

GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. On the Properties File page, configure the required properties as needed. Look for the ones marked as #TODO. And then click **Next**. Some properties to consider modifying include:
 - a. `gg.eventhandler.s3.region`: name of the region where Amazon S3 bucket is located. Please make sure that Amazon S3 bucket and Amazon Redshift cluster are in the same region.
 - b. `gg.eventhandler.s3.bucketMappingTemplate`: Amazon S3 bucket name that will be used as staging area.
6. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Amazon S3

Learn to create a Replicat process for an Amazon S3 target in OCI GoldenGate

Before you begin

Before adding and running a Replicat for Amazon S3, ensure that you have the following:

- An Amazon S3 connection created and assigned to your target Big Data deployment
- An Amazon Web Services (AWS) region

Add a Replicat for Amazon S3

Learn to add a replicat process for Amazon S3.

To add a Replicat for an Amazon S3 target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The types of Replicats are:

- Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a short description to distinguish this process from others.
 - c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **Amazon S3** from the dropdown.
 - e. For **Available Aliases for Amazon S3**, select your alias from the dropdown.
 4. On the Parameter File page, add and configure Replicat parameters as needed, and then click **Next**.

 **Note:**

- If using Coordinated Replicat, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/
Deployment/etc/conf/ogg/<replicat
name>.properties
```

Using coordinated replicat will result in multiple files created.

If omitted, you will encounter the following error:

```
OGG-01091 Oracle GoldenGate Delivery, RSNOW.prm: Unable to
open file
"/u02/Deployment/etc/conf/ogg/<replicat
name>001.properties" (error 2, No such file or
directory).
```

- GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. On the Properties File page, configure the following Amazon S3 properties:
 - `gg.eventhandler.s3.region` provide the AWS region for the target S3 bucket.
 - `gg.eventhandler.s3.bucketMappingTemplate`: provide the target S3 bucket name. If the bucket doesn't exist, it can but auto created by OCI GoldenGate. You can provide static bucket names, or use Template Keywords to assign bucket names dynamically.

 **Note:**

Ensure the bucket name contains only lowercase characters. Uppercase characters can cause the Replicat to fail.

6. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Kafka

Learn to add a replicat process for Kafka.

To add a Replicat for an Apache Kafka target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The types of Replicats are:

- Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a short description to distinguish this process from others.
 - c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **Kafka** from the dropdown.
 - e. For **Available aliases for OCI**, select the Kafka connection.
 4. On the Parameter File page, add and configure Replicat parameters as needed, and then click **Next**.

 **Note:**

GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")  
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. On the Properties File page, for `gg.handler.kafkahandler.topicMappingTemplate=`, enter the stream name.
6. Click **Add and Run**.

You're returned to the Administration Service Overview page where you can monitor the status of the Replicat process. Click the process name to view its details and access reports.

Add a Replicat for Confluent Kafka

Learn to add a replicat process for Confluent Kafka.

To add a Replicat for an Confluent Kafka target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The types of Replicats are:

- Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.
 - c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **Kafka** from the dropdown.
 - e. For **Available Alias**, select the Confluent Kafka connection.
 - f. Enable **Kafka Connect**.
 - g. For **Converter**, select **Avro**.
 - h. For **Schema Registry**, select the Confluent Schema Registry connection.
 4. On the Parameter File page, add and configure Replicat parameters as needed, and then click **Next**.

Note:

GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")  
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. On the Properties File page, configure the following Kafka Connect property:
 - For `gg.handler.kafkaconnect.topicMappingTemplate`: provide the Confluent Kafka topic name.
6. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Elasticsearch Server

Learn to add a replicat process for Elasticsearch Server.

To add a Replicat for an Elasticsearch Server target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.
The types of Replicats are:
 - Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. If the source is Trail, then for **Trail Name**, enter the name of the Extract trail.
 - c. For **Target**, select **Elasticsearch** from the dropdown.
 - d. For **Alias**, select your Elasticsearch connection.
4. On the Parameter File page, leave the default, and then click **Next**.
5. On the Properties File page, leave the field blank. For additional Elasticsearch Server replication properties, you can refer to Configuring the Elasticsearch Handler.
6. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Google BigQuery

Learn to create a Replicat process for a Google Cloud Storage target.

Before you begin

Before you create a Replicat for Google BigQuery, ensure that you:

- Create and assign both a Google Cloud Storage and a Google BigQuery connection to the Big Data deployment. The deployment uses Google Cloud Storage as a staging location for Google BigQuery.

 **Note:**

Ensure that the Google Cloud Storage bucket and BigQuery dataset exist in the same location or region.

- Have your [Google Cloud Service Account Key](#).
- Assign the appropriate Google BigQuery permissions

Add a Replicat

To add a Replicat for an Google BigQuery target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).

2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The types of Replicats are:

- Classic Replicat
- Coordinated Replicat

3. On the Replicat Options page, complete the following fields, and then click **Next**:

- a. For **Process Name**, enter a name for the Replicat process.

- b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.

- c. For **Trail Name**, enter two-character trail name.

- d. For **Target**, select **Google BigQuery** from the dropdown.

- e. For **Available Aliases**, select your alias from the dropdown.

- f. For **Available staging locations**, Google Cloud Storage should already be selected.

- g. For **via staging alias**, select a Google Cloud Storage connection.

4. On the Parameter File page, add and configure Replicat parameters as needed, and then click **Next**.

 **Note:**

- If using Coordinated Replicat, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/
Deployment/etc/conf/ogg/<replicat
name>.properties
```

Using coordinated replicat will result in multiple files created.

If omitted, you will encounter the following error:

```
OGG-01091 Oracle GoldenGate Delivery, RSNOW.prm: Unable
to open file
"/u02/Deployment/etc/conf/ogg/<replicat
name>001.properties" (error 2, No such file or
directory).
```

- GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. On the Properties File page, locate `gg.eventhandler.gcs.bucketMappingTemplate=<gcs bucket>`, and then replace `<gcs bucket>` with the Google Cloud Storage bucket name.
6. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Google Cloud Storage

Learn the prerequisites and process to add a Replicat for Google Cloud Storage.

Before you begin

Before you create a Replicat for Google Cloud Storage, ensure that you have:

- Your Google Cloud Storage account and [Google Cloud Service Account Key](#).
- Google Cloud Storage bucket and object permissions

Add a Replicat

Learn to add a replicat process for Google Cloud Storage.

To add a Replicat for an Google Cloud Storage target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.
The types of Replicats are:
 - Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.
 - c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **Google Cloud Storage** from the dropdown.
 - e. For **Available Aliases**, select your alias from the dropdown.
4. On the Parameter File page, add and configure Replicat parameters as needed, and then click **Next**.

 **Note:**

- If using Coordinated Replicat, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/
Deployment/etc/conf/ogg/<replicat
name>.properties
```

Using coordinated replicat will result in multiple files created.

If omitted, you will encounter the following error:

```
OGG-01091 Oracle GoldenGate Delivery, RSNOW.prm: Unable to
open file
"/u02/Deployment/etc/conf/ogg/<replicat
name>001.properties" (error 2, No such file or
directory).
```

- GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. On the Properties File page, locate `gg.eventhandler.gcs.bucketMappingTemplate=<gcs bucket>`, and then replace `<gcs bucket>` with the Google Cloud Storage bucket name.

6. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Azure Cosmos DB for MongoDB

Learn to add a Replicat process for an Azure Cosmos DB for MongoDB target.

To add a Replicat for an Azure Cosmos DB for MongoDB target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.
The types of Replicats are:
 - Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.
 - c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **Azure Cosmos DB for Mongo DB** from the dropdown.
 - e. For **Available aliases for OCI**, select the assigned Azure Cosmos DB for MongoDB connection.
4. On the Parameter File page, add and configure Replicat parameters as needed, such as `MAP *.*`, `TARGET *.*`; and then click **Next**.

 **Note:**

- If using Coordinated Replicat, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/
Deployment/etc/conf/ogg/<replicat
name>.properties
```

Using coordinated replicat will result in multiple files created.

If omitted, you will encounter the following error:

```
OGG-01091 Oracle GoldenGate Delivery, RSNOW.prm: Unable to
open file
"/u02/Deployment/etc/conf/ogg/<replicat
name>001.properties" (error 2, No such file or
directory).
```

- GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Azure Data Lake Storage

Learn to add a replicat process for Azure Data Lake Storage.

To add a Replicat for an Azure Data Lake Storage target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The types of Replicats are:

- Classic Replicat
- Coordinated Replicat

3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.

- c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **Azure Data Lake Storage** from the dropdown.
 - e. For **Alias**, select the Azure Data Lake Storage connection.
4. On the Parameter File page, add and configure Replicat parameters as needed, and then click **Next**.

 **Note:**

- If using Coordinated Replicat, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/  
Deployment/etc/conf/ogg/<replicat  
name>.properties
```

Using coordinated replicat will result in multiple files created.

If omitted, you will encounter the following error:

```
OGG-01091 Oracle GoldenGate Delivery, RSNOW.prm: Unable  
to open file  
"/u02/Deployment/etc/conf/ogg/<replicat  
name>001.properties" (error 2, No such file or  
directory).
```

- GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")  
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. On the Properties File page, configure the following Azure Data Lake Storage properties:
- Required property:
- For `gg.eventhandler.abs.bucketMappingTemplate`: provide the Azure Data Lake Storage container name. If container is pre-configured, a static container name can be provided. If Azure authentication method permissions are provided, then OCI GoldenGate can use Template Keywords to auto create the container.
- (Optional) Additional properties you may consider adding:
- `gg.handler.abs.format`: Select how to format the output. JSON is the default setting. Available options include:
 - delimitedtext
 - json
 - json_row

- xml
- avro_row_ocf
- avro_op_ocf

6. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Azure Event Hubs

Learn to add a replicat process for Azure Event Hubs.

To add a Replicat for an Azure Event Hubs target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select **Classic Replicat** and then click **Next**.
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.
 - c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **Azure Event Hubs** from the dropdown.
 - e. For **Alias**, select the Azure Event Hubs connection.
4. On the Replicat Parameters page, leave the default, and then click **Next**.
5. On the Properties page, provide a topic name in `topicMappingTemplate`. Topic name can be a static name or a Template Keywords for a dynamic topic name. If topic is not present, it will be auto-created by OCI GoldenGate.
6. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Azure Synapse Analytics

Learn to add a replicat process for Azure Synapse Analytics.

To add a Replicat for an Azure Synapse Analytics target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a **Replicat type**, and then click **Next**.

The types of Replicats are:

- Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:

- a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a short description to distinguish this process from others.
 - c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **Azure Synapse Analytics** from the dropdown.
 - e. For **Alias**, select the Azure Data Lake Storage connection.
4. On the Parameter File page, add and configure Replicat parameters as needed, and then click **Next**.

 **Note:**

- If using Coordinated Replicat, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/
Deployment/etc/conf/ogg/<replicat
name>.properties
```

If omitted, you will encounter the following error:

```
OGG-01091 Oracle GoldenGate Delivery, RSNOW.prm: Unable
to open file
"/u02/Deployment/etc/conf/ogg/<replicat
name>001.properties" (error 2, No such file or
directory).
```

- GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. On the Properties File page, configure the following Azure Synapse Analytics for
- For `gg.eventhandler.abs.bucketMappingTemplate`: provide the Azure Storage Container associated with the Azure Synapse Analytics workspace.
 - For `gg.eventhandler.synapse.credential`: provide the name of the credential used to authenticate the Azure Storage Container associated with the Azure Synapse Analytics workspace.
6. Click **Add and Run**.

You return to the Administration Service Overview page where you can monitor the status of the Replicat process. Click the process name to view its details and access reports.

Add a Replicat for MongoDB

Learn to add a Replicat process for a MongoDB target.

To add a Replicat for a MongoDB target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The types of Replicats are:

- Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. (Optional) For **Description**, enter a brief description to distinguish this Replicat from others.
 - c. For **Trail Name**, enter two-character trail name.
 - d. For **Target**, select **MongoDB** from the dropdown.
 - e. For **Available Aliases**, select your alias from the dropdown.
 4. On the Parameter File page, add and configure Replicat parameters as needed, and then click **Next**.

 **Note:**

- If using Coordinated Replicat, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/
Deployment/etc/conf/ogg/<replicat
name>.properties
```

Using coordinated replicat will result in multiple files created.

If omitted, you will encounter the following error:

```
OGG-01091 Oracle GoldenGate Delivery, RSNOW.prm: Unable
to open file
"/u02/Deployment/etc/conf/ogg/<replicat
name>001.properties" (error 2, No such file or
directory).
```

- GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat for Redis

Learn to add a replicat process for Redis.

To add a Replicat for an Redis target:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.

The types of Replicats are:

- Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. If the source is Trail, then for **Trail Name**, enter the name of the Extract trail.
 - c. For **Target**, select **Redis**.

- d. For **Alias**, select your Redis connection.
4. On the Parameter File page, add and configure Replicat parameters as needed, and then click **Next**.

 **Note:**

When using Coordinated Replicat, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/Deployment/etc/  
conf/ogg/<replicat name>.properties
```

5. On the Properties page, you can update Redis integration point by changing the `gg.handler.redis.integrationType` property. OCI GoldenGate supports *hashmaps*, *streams* and *jsons* as Redis integration points. Refer to Redis Handler Configuration Properties for more information.
6. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Add a Replicat

To add a Replicat for Snowflake:

1. In the OCI GoldenGate deployment console, on the Administration Service Overview page, click **Add Replicat** (plus icon).
2. On the Add Replicat page, select a Replicat type, and then click **Next**.
The types of Replicats are:
 - Classic Replicat
 - Coordinated Replicat
3. On the Replicat Options page, complete the following fields, and then click **Next**:
 - a. For **Process Name**, enter a name for the Replicat process.
 - b. For **Trail Name**, enter two-character trail name.
 - c. For **Target**, select the target **Snowflake** connection from the dropdown.
 - d. For **Available Aliases**, select your alias from the dropdown.
 - e. (Optional) **Enable external storage** to select an available staging location from the dropdown.
4. On the Replicat Parameters page, add the necessary mappings, and then click **Next**:

 **Note:**

- If using Coordinated Replicat, add the following parameters on the second line of the Parameter File:

```
TARGETDB LIBFILE libggjava.so SET property=/u02/
Deployment/etc/conf/ogg/<replicat
name>.properties
```

Using coordinated replicat will result in multiple files created.

If omitted, you will encounter the following error:

```
OGG-01091 Oracle GoldenGate Delivery, RSNOW.prm: Unable
to open file
"/u02/Deployment/etc/conf/ogg/<replicat
name>001.properties" (error 2, No such file or
directory).
```

- GoldenGate uses Greenwich Mean Time (GMT) by default. Use SETENV to override the default setting. For example:

```
setenv (TZ="US/Eastern")
setenv (TZ="GMT+5")
```

Learn more about SETENV.

5. (Optional) On the Properties File page, review the following properties:

 **Note:**

These properties are set when you created the connection and shouldn't be modified here.

- `ggs.eventhandler.snowflake.connectionURL`: JDBC URL to connect to Snowflake.
 - `ggs.eventhandler.snowflake.UserName`: Snowflake database username.
 - `ggs.eventhandler.snowflake.Password`: Password associated with the Snowflake database user.
 - `ggs.eventhandler.snowflake.storageIntegration`: The credential for Snowflake data warehouse to access the respective Object store files. For more information, see Snowflake storage integration.
6. Click **Create and Run**. If you click **Create**, then you can manually start the Replicat later from the Administration Service Overview page.

You return to the Administration Service Overview page where you can view the creation of the Replicat process and access Replicat Actions.

Using the Admin Client

Admin Client is a command line utility for controlling and configuring tasks in Oracle GoldenGate and OCI GoldenGate.

Access Admin Client

Use Admin Client to connect to OCI GoldenGate to configure tasks and view process information and log messages. You can launch Admin Client one of two ways:

- Click **Launch Admin Client** on the deployment details page.
- Launch **Cloud Shell**, and then run Admin Client.

Note:

If you have an Oracle GoldenGate Marketplace version running on a Compute instance, you can [access its AdminClient to connect to your OCI GoldenGate deployment](#).

Connect to Admin Client through Cloud Shell

Run the following commands to connect to a OCI GoldenGate deployment that has a public endpoint in Cloud Shell:

```
> adminclient
OGG (not connected) 1> connect <deployment-public-url-or-ip> as <goldengate-
user> password <goldengate-password> !
```

However, connecting to a private OCI GoldenGate requires creating a bastion, bastion session, and SSH tunnel.

To connect to a private OCI GoldenGate deployment in Admin Client:

1. In the Oracle Cloud console global navigation bar, click **Cloud Shell**. If this is your first time connecting to Cloud Shell, it will take a few moments to connect.
2. You can run the following command to generate SSH keys, or skip this step and generate the keys when you [create the bastion](#):

```
ssh-keygen -t rsa
```

Keep the default filename and don't enter a passphrase when prompted. The private key is located at `~/ssh/id_rsa` and the public key is located at `~/ssh/id_rsa.pub`.

3. On the deployment details page, take note of the deployment's **Private IP** and **Subnet** information.
4. Create a Bastion.
 - a. From the Oracle Cloud console menu, select **Identity & Security**, and then select **Bastion**.

- b. Click **Create Bastion**.
 - c. In the Create Bastion panel, enter a name, and then select the same subnet where the deployment resides.
 - d. For CIDR block allowlist, enter `0.0.0.0/0`.
 - e. Click **Create bastion**.
5. Create a session.
 - a. After the bastion is in an Active state, on the bastion details page, click **Create Session**.
 - b. For Session type, **select SSH Port forwarding session**.
 - c. Enter a name for the session.
 - d. For Connect to target using, select IP Address, and then enter the deployment's private IP.
 - e. For Port, enter 443.
 - f. For Add SSH Key, copy and paste the contents of the public key (`~/ssh/id_rsa.pub`) from Cloud Shell.
 - g. Click **Create Session**.
6. After the bastion session state is active, select **View SSH command** from its Action menu (ellipsis icon).
7. In the View SSH command dialog, enter the path to the private key (`~/ssh/id_rsa`) in place of `<private-key>` and replace `<localport>` with the port in Cloud Shell that will forward the connection to the bastion.

 **Note:**

Cloud Shell doesn't allow port forwarding on a privileged port with sudo access, so you must use a non-privileged port like 7443. After the command runs once in the foreground to add the Bastion host to `known_hosts`, you can append an ampersand (`&`) to the end of the command so that it can run in the background next time.

8. Copy the command and then run it in Cloud Shell. You can ignore `bind: Cannot assign requested address` messages.
9. Start the admin client.

```
adminclient
```
10. Connect to the OCI GoldenGate deployment.

```
connect 127.0.0.1:7443 as <goldengate-user> password <goldengate-  
password> !
```

 **Note:**

The exclamation point (!) at the end of the command is very important. Without it, the command fails and returns an error.

Use Admin Client

After connecting successfully, you can run any of the following commands:

- Display the status of OCI GoldenGate processes:

```
info all
```

- View statistics for your Extract:

```
view stats
```

- View the content of the ggserror log file:

```
view messages
```

- Purge trail files that are no longer used by Extracts:

```
purge exttrail <trail-file-name>
```

See AdminClient Command Line Interface Commands for the full list of commands.

6

Stream and analyze

Learn to stream, enrich, and analyze data in OCI GoldenGate Stream Analytics.

Articles in this section:

- [About Stream Analytics](#)
- [Create Stream Analytics resources](#)
- [Access the Stream Analytics deployment](#)
- [Use Stream Analytics](#)

About Stream Analytics

Create custom operational dashboards that provide real time monitoring and analyses of event streams using OCI GoldenGate Stream analytics. Identify events of interest, run queries against the event streams in real time, or raise alerts based on your analysis.



Note:

Stream Analytics is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.

Stream analytics concepts

The following concepts are essential for working with OCI GoldenGate Stream Analytics:

- **Connection:** Stores the connectivity information for a source or target technology.
- **Stream:** A continuous flow of dynamic data.
- **Pipeline:** The workflow data from source to target.
- **Business logic:** Various filters and functions you can apply to a pipeline to obtain the precise data you want to analyze.
- **Publishing:** Makes the pipeline available to all Stream analytics users and sends data to targets.

Stream analytics limitations

While OCI GoldenGate Stream Analytics appears the same as Oracle Stream Analytics, there are certain features that are not supported in the OCI version. Pay careful attention to notes in the Oracle Stream Analytics that inform you of whether or not a feature is supported in OCI GoldenGate Stream Analytics.

Supported connections

Learn about what types of connections are supported by OCI GoldenGate Stream Analytics.

OCI GoldenGate Stream Analytics supports the following **source** technology types:

- Oracle Autonomous Database
- Oracle Database
- OCI Streaming
- OCI MySQL Database Service
- Apache Kafka
- Confluent Kafka
- Oracle GoldenGate server

 **Note:**

You can also create Coherence, Ignite, and Java Message Server (JMS) connections directly within the Stream Analytics console.

Stream Analytics supports the following **target** technology types:

- Oracle Autonomous Database
- Oracle Database
- OCI Object Storage
- Apache Kafka
- OCI Streaming

 **Note:**

You can also create Amazon S3, Azure Data Lake Storage, Coherence, Hadoop File Storage (HDFS), Ignite, JMS, and MongoDB connections directly within the Stream Analytics console.

Metering and billing for Stream Analytics deployments

Ensure that you review the information in Metering and billing for OCI GoldenGate deployments about Oracle Compute Unit (OCPU) selection and scaling.

OCI GoldenGate Stream Analytics OCPU usage is calculated based on the following factors:

- Stream Analytics console
- Number of Streaming pipelines
- Ignite cluster

- GoldenGate Big Data cluster

Before calculating the number of OCPUs you need, let's first review how many compute units each Stream Analytics resource requires. 1 OCPU is equal to 2 compute units (vCPUs). 1 vCPU is equal to 1000 millicores (1000m).

The following table lists example Stream Analytics pipeline settings and the calculated number of OCPUs required.

Pipeline	Driver	Executor	Total vCPUs	OCPUs billed
Pipeline A	500m	1 x 500m	1000m	1
Pipeline B	500m	2 x 500m	1500m	1
Pipeline C	500m	4 x 500m	2500m	2
Pipeline D	600m	2 x 700m	2000m	1
Pipeline E	1000m	2 x 1000m	3000m	2

You can configure the Driver and Executor settings as needed for each pipeline in the Stream Analytics console.

The following table lists example Stream Analytics resource configurations based on the number of pipelines (from the above table) and the calculated number of OCPUs required.

Stream Analytics console	Number of pipelines	Streaming pipelines	Ignite cluster	GoldenGate for Big Data cluster	OCPUs billed
1000m	1 x Pipeline A	1000m	0	0	1
1000m	3 x Pipeline A	3000m	0	0	2
1000m	1 x Pipeline B	1500m	0	0	2
1000m	1 x Pipeline B	1500m	2 x 500m	500m	2
1000m	1 x Pipeline A 1 x Pipeline B	2500m	2 x 500m	500m	3
1000m	2 x Pipeline A 1 x Pipeline B	3500m	2 x 500m	500m	3

The Stream Analytics console requires 1000m. Each streaming pipeline requires additional millicores depending on their settings. The Ignite cluster, if activated, requires a minimum of 2 cluster instances. You can configure the millicore limit for both Ignite and GoldenGate Big Data clusters in the Stream Analytics console. When added together, you can determine the total number of OCPUs that you need to select when creating your Stream Analytics deployment.

If you're unsure, you can start with 2 or 3 OCPUs, and then review the OCPU consumption metrics on the deployment details page and adjust accordingly.

Create Stream Analytics resources

To use Stream Analytics in OCI GoldenGate, create a Stream Analytics deployment and assign connections.

 **Note:**

Stream Analytics is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.

 **Tip:**

Ensure that you review the OCPU management and billing information for both OCI GoldenGate and Stream Analytics deployments before you proceed.

Create the Stream Analytics deployment

To create a Stream Analytics deployment in the Oracle Cloud console:

1. In the Console navigation menu, click **Oracle Database**, and then select **GoldenGate**.
2. On the Deployments page, click **Create deployment**.
3. In the Create deployment panel, enter a name and optionally, a description.
4. From the Compartment dropdown, select a compartment in which to create the deployment.
5. For **OCPU count** enter the number of Oracle Compute units (OCPU) to use.

 **Note:**

One OCPU is equivalent to 16GB of memory. 3 to 4 OCPUs is sufficient for one Stream analytics pipeline. For more information, see OCPU management and billing.

6. (Optional) Select **Auto scaling**.

 **Note:**

Auto scaling enables OCI GoldenGate to scale up to three times the number of OCPUs you specify for OCPU Count, up to 24 OCPUs. For example, if you specify your OCPU Count as 2 and enable Auto Scaling, then your deployment can scale up to 6 OCPUs. If you specify your OCPU Count as 20 and enable Auto Scaling, OCI GoldenGate can only scale up to 24 OCPUs.

7. From the **Subnet in <Compartment>** dropdown, select the subnet to which a private endpoint is created from the OCI GoldenGate service tenancy. This ensures that the deployment is always available over this subnet, as long as the policies for this subnet allow access. The private endpoint is only used to access

the deployment console, and doesn't provide access to other resources in the subnet. To select a subnet in a different compartment, click **Change compartment**.

 **Note:**

You can only select a private subnet when creating a deployment.

8. Select a license type.
9. (Optional) Click **Show advanced options** for network options and to add tags.
 - a. In the **Network** tab,
 - i. Select **Enable GoldenGate console public access** to include a public endpoint in addition to a private endpoint, and allow public access to the deployment console for users. If selected, OCI GoldenGate creates a load balancer in your tenancy to create a public IP. Select a subnet in the same VCN as this deployment in which to create the load balancer.

 **Note:**

The load balancer is a resource that comes with an additional cost. You can manage this resource, but ensure that you don't delete the load balancer while your deployment is still in use. [Learn more about load balancer pricing](#).

- ii. Select **Customize endpoint** to provide a private fully qualified domain name (FQDN) prefix that you'll use to access the private service console URL. You can also optionally upload an SSL/TLS certificate (.pem) and its corresponding private key, however, password protected certificates are not supported. It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.
A self-signed certificate is generated for you, if you don't provide one.

 **Note:**

It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.

- b. In the **Maintenance** tab:
 - i. Select **Customize maintenance window** to define the start of the maintenance window to upgrade the deployment.
 - ii. (Optional) For **Major release auto-upgrade period in days**, enter the number of days, between 0 and 365.
 - iii. (Optional) For **Bundle release auto-upgrade period in days**, enter the number of days, between 0 and 180 days.
 - iv. (Optional) For **Security patch auto-upgrade period in days**, enter the number of days, between 0 and 14 days.
 - v. Select **Enable interim release auto-upgrade**, and, optionally, enter the number of days.

 **Note:**

Learn more about scheduling upgrades.

- c. In the **Tags** tab, add tags to help track the resources within your tenancy. Click **+ Additional tag** to add more tags. [Learn more about tagging.](#)
10. Click **Next**.
11. Select **Stream analytics** for deployment type.
12. The Stream analytics technology type is automatically selected for you.
13. For **GoldenGate instance name**, enter a name for the stream analytics instance.
14. For **Administrator username**, enter `osaadmin`.
15. For **Administrator password**, enter a password, and then confirm that password.
16. Click **Create**.

Your Stream analytics deployment takes a few minutes to create. Its status changes to Active when your deployment is ready to use. Ensure that you create and assign connections to the deployment before you use your deployment.

Supported connections

Learn about what types of connections are supported by OCI GoldenGate Stream Analytics.

OCI GoldenGate Stream Analytics supports the following **source** technology types:

- Oracle Autonomous Database
- Oracle Database
- OCI Streaming
- OCI MySQL Database Service
- Apache Kafka
- Confluent Kafka
- Oracle GoldenGate server

 **Note:**

You can also create Coherence, Ignite, and Java Message Server (JMS) connections directly within the Stream Analytics console.

Stream Analytics supports the following **target** technology types:

- Oracle Autonomous Database
- Oracle Database
- OCI Object Storage
- Apache Kafka

- OCI Streaming

 **Note:**

You can also create Amazon S3, Azure Data Lake Storage, Coherence, Hadoop File Storage (HDFS), Ignite, JMS, and MongoDB connections directly within the Stream Analytics console.

Assign a connection to a deployment

Ensure that you have connections created for your source and target technologies.

To assign a connection to a deployment:

1. On the deployment details page, under **Resources**, click **Assigned connections**.
2. Click **Assign connection**.
3. In the Assign connection dialog, select a connection from the dropdown. If you want to select a connection from a different compartment, click **Change Compartment**.
4. Click **Assign connection**.

The selected connection appears in the Assigned connections list. You can also view and manage this relationship from the Connection details page under **Assigned deployments**.

Access the Stream Analytics deployment

After you create the Stream Analytics deployment and assign connections, you can access the Stream Analytics console from the deployment details page.

 **Note:**

Stream Analytics is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.

To access the Stream Analytics deployment:

1. On the OCI GoldenGate Deployments page, select the Stream Analytics deployment.
2. On the Stream Analytics deployment details page, click **Launch console**.
Alternatively, you can copy the **Console url** and paste it into your browser.
3. On the Stream Analytics log in page, enter the Administrator username and Administrator password provided when you created the deployment.

You can now use Stream Analytics to create streams and pipelines. See Use Stream Analytics for next steps.

Use Stream Analytics

If this is your first time using OCI GoldenGate Stream Analytics, use this information to get yourself started. If you already created your Stream Analytics deployment, you can skip directly to Launch the deployment console.

Stream and analyze taskflow



Note:

Stream Analytics is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.

Task	Description	More information
Review Security best practices	Develop a firm understanding of your responsibilities to keep your OCI GoldenGate deployments and connections secure.	Securing OCI GoldenGate
Create OCI resources	This task is typically performed by an administrator. Ensure that the required networking resources are created before you begin.	Create Oracle Cloud resources
Create deployments	A deployment is a container for your OCI GoldenGate resources.	Example OCI GoldenGate topologies Create deployments
Create connections	A connection contains the network connectivity details for a data source or target.	About connections
Assign connections to deployments	To use a connection as a source or target, you must assign it to a deployment.	Create an association between connections and deployments
Launch the Stream Analytics console	Create your data replication processes in the Stream Analytics console.	Access the deployment
Create a stream	A Stream is a source of continuous and dynamic data. The data can be from a wide variety of data sources such as IoT sensors to information from geospatial services or social networks.	Create Streams
Create a pipeline	A pipeline includes a sequence of data processing stages such as, Query, Pattern, Rule, Query Group, Custom, and Scoring.	Create a Pipeline

Task	Description	More information
Add Business Logic	Transform the input data stream, add business logic to the pipeline to analyze the input data stream.	See Transform See Analyze
Publish the Pipeline	To make the pipeline available for all the users of Oracle Stream Analytics and send data to targets, you must publish a pipeline.	Publishing a Pipeline

7

Transform data

Learn to design data transformations using OCI GoldenGate Data Transforms.

Articles in this section:

- [About Data Transforms](#)
- [Create Data Transforms resources](#)
- [Access the Data Transforms deployment](#)
- [Use Data Transforms](#)

About Data Transforms

Design graphical data transformations such as data flows and workflows to move and transform data between systems.



Note:

Data Transforms is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.

Data Transforms concepts

Whether you're new to Data Transforms or have past experience with Oracle Data Integrator, it's helpful to familiarize yourself with these concepts before starting with OCI GoldenGate Data Transforms. See Terminology Information in the *Using Data Transforms* guide to learn more.

Supported connection types for Data Transforms

See Supported Connection Types for a full list of connections that you can use with OCI GoldenGate Data Transforms.



Tip:

For each connection type used with OCI GoldenGate Data Transforms, you must create a Generic connection, and assign the connection to the Data Transforms deployment.

After assigning the Generic connection(s) to the Data Transforms deployment, launch the console from the deployment details page and log in. In the Data Transforms deployment

console, create a connection for each data source using the Host name(s) provided for the Generic connection(s).

Create Data Transforms resources

To use Data Transforms in OCI GoldenGate, create a Data Transforms deployment, and then create and assign a Generic connection type to the deployment.

Note:

Data Transforms is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.

Create a Data Transforms deployment

To create a Data Transforms deployment:

1. In the Console navigation menu, click **Oracle Database**, and then select **GoldenGate**.
2. On the Deployments page, click **Create deployment**.
3. In the Create deployment panel, enter a name and optionally, a description.
4. From the Compartment dropdown, select a compartment in which to create the deployment.
5. Select one of the following options:
 - **Production**: Sets up a deployment with recommended defaults for a production environment. The minimum number of OCPUs is 4, with auto-scaling enabled.
 - **Development or testing**: Sets up a deployment with recommended defaults for a development or testing environment. The minimum number of OCPUs is 1.
6. For **OCPU count** enter the number of Oracle Compute units (OCPUs) to use.

Note:

One OCPU is equivalent to 16gb of memory. For more information, see OCPU management and billing.

7. (Optional) Select **Auto scaling**.

 **Note:**

Auto scaling enables OCI GoldenGate to scale up to three times the number of OCPUs you specify for OCPU Count, up to 24 OCPUs. For example, if you specify your OCPU Count as 2 and enable Auto Scaling, then your deployment can scale up to 6 OCPUs. If you specify your OCPU Count as 20 and enable Auto Scaling, OCI GoldenGate can only scale up to 24 OCPUs.

8. From the **Subnet in <Compartment>** dropdown, select the subnet to which a private endpoint is created from the OCI GoldenGate service tenancy. This ensures that the deployment is always available over this subnet, as long as the policies for this subnet allow access. The private endpoint is only used to access the deployment console, and doesn't provide access to other resources in the subnet.

To select a subnet in a different compartment, click **Change compartment**.

 **Note:**

You can only select a private subnet when creating a deployment.

9. Select a license type.
10. (Optional) Click **Show advanced options** for network options and to add tags.
 - a. In the **Network** tab,
 - i. Select **Enable GoldenGate console public access** to include a public endpoint in addition to a private endpoint, and allow public access to the deployment console for users. If selected, OCI GoldenGate creates a load balancer in your tenancy to create a public IP. Select a subnet in the same VCN as this deployment in which to create the load balancer.

 **Note:**

The load balancer is a resource that comes with an additional cost. You can manage this resource, but ensure that you don't delete the load balancer while your deployment is still in use. [Learn more about load balancer pricing](#).

- ii. Select **Customize endpoint** to provide a private fully qualified domain name (FQDN) prefix that you'll use to access the private service console URL. You can also optionally upload an SSL/TLS certificate (.pem) and its corresponding private key, however, password protected certificates are not supported. It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.
A self-signed certificate is generated for you, if you don't provide one.

 **Note:**

It is your responsibility to ensure the FQDN is resolvable in the subnet you have previously selected.

- b. In the **Maintenance** tab:
 - i. Select **Customize maintenance window** to define the start of the maintenance window to upgrade the deployment.
 - ii. (Optional) For **Major release auto-upgrade period in days**, enter the number of days, between 0 and 365.
 - iii. (Optional) For **Bundle release auto-upgrade period in days**, enter the number of days, between 0 and 180 days.
 - iv. (Optional) For **Security patch auto-upgrade period in days**, enter the number of days, between 0 and 14 days.
 - v. Select **Enable interim release auto-upgrade**, and, optionally, enter the number of days.

 **Note:**

Learn more about scheduling upgrades.

- c. In the **Tags** tab, add tags to help track the resources within your tenancy. Click **+ Additional tag** to add more tags. [Learn more about tagging.](#)
11. Click **Next**.
 12. On the GoldenGate details page, for **Choose a deployment type**, select **Data transforms**.
 13. For Version, the latest version is automatically selected. Click **Change version** to select a different version.

 **Note:**

Learn more about versions.

14. For **GoldenGate instance name**, enter the name that the deployment will assign to the GoldenGate deployment instance upon creation.
15. The Administrator username automatically populates with SUPERVISOR for you.
16. For Administrator password, enter a password for the SUPERVISOR user, and then confirm that password.
17. Click **Create**.

The deployment takes a few minutes to create. Its status changes to Active when it's ready to use. You can then click Launch console from the deployment details page, or select Launch console from the deployment's Actions (three dots) menu on the Deployments page.

Assign a connection to a deployment

Ensure that you have connections created for your source and target technologies.

To assign a connection to a deployment:

1. On the deployment details page, under **Resources**, click **Assigned connections**.

2. Click **Assign connection**.
3. In the Assign connection dialog, select a connection from the dropdown. If you want to select a connection from a different compartment, click **Change Compartment**.
4. Click **Assign connection**.

The selected connection appears in the Assigned connections list. You can also view and manage this relationship from the Connection details page under **Assigned deployments**.

Access the Data Transforms deployment

After you create the Data Transforms deployment and assign connections, you can access the Data Transforms console from the deployment details page.

Note:

Data Transforms is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.

To access the Data Transforms deployment:

1. On the OCI GoldenGate Deployments page, select the Data Transforms deployment.
2. On the Data Transforms deployment details page, click **Launch console**.
Alternatively, you can copy the **Console URL** and paste it into your browser.
3. On the Data Transforms log in page, enter the Administrator username (`SUPERVISOR`) and password provided when you created the deployment.

You can now use Data Transforms to create streams and pipelines. See Use Data Transforms for next steps.

Use Data Transforms

If this is your first time using OCI GoldenGate Data Transforms, use this information to get yourself started. If you already created your Data Transforms deployment, you can skip directly to Launch the Data Transforms deployment console.

Data Transforms taskflow

Note:

Data Transforms is currently in Limited Availability and only available in specific regions. To learn more and gain access, contact your Oracle representative or Oracle Support.

Task	Description	More information
Review Security best practices	Develop a firm understanding of your responsibilities to keep your OCI GoldenGate deployments and connections secure.	Securing OCI GoldenGate
Create OCI resources	This task is typically performed by an administrator. Ensure that the required networking resources are created before you begin.	Create Oracle Cloud resources
Create deployments	A deployment is a container for your OCI GoldenGate resources.	Create a Data Transforms deployment
Create Generic connections	Create a generic connection for each Data Transforms data source.	Create a Generic connection
Assign connections to deployments	To use a connection as a source or target, you must assign it to a deployment.	Assign a connection to a deployment
Launch the Data Transforms console	Create your data flows and workflows in Data Transforms.	Access Data Transforms
Create connections	Create connections in Data Transforms to data sources to use in a project.	Work with connections
Create a project	A project is the top-level container, which can include multiple folders to organize your data flows or work flows into logical groups.	Work with Projects
Create and Run a Data Load	A data load allows you to load multiple data entities from a source connection to a target connection.	Create a Data Load Run a Data Load
Monitor Status of Data Loads, Data Flows, and Workflows	When you run a data load, data flow, or workflow Oracle Data Transforms runs jobs in the background to complete the request. You can view the status of the job in the panel on the bottom right of the page.	Monitor Status of Data Loads, Data Flows, and Workflows
Create Data Flows and Workflows	A data flow defines how the data is moved and transformed between different systems. A workflow is made up of multiple flows organized in a sequence in which they must be executed.	Create a Data Flow Create a New Workflow

8

Manage

Learn to manage and monitor your resources to keep your deployments and replication processes running smoothly.

Articles in this section:

- [Manage connections](#)
- [Manage deployments](#)
- [Manage deployment users](#)
- [Manage Trail files](#)
- [Manage master encryption key wallets](#)
- [Manage Truststore certificates](#)
- [Manage deployment backups](#)
- [Monitor performance](#)

Manage connections

Learn connection management tasks including how to edit, unassign, move, and delete connections.

View connection details

Select a connection from the Connections list to view its details. On the Connection details page, you can:

- View the connection's status, which can be one of the following:
 - Creating
 - Updating
 - Active
 - Deleting
 - Deleted
 - Failed
- Perform actions on the connection such as, edit, move, add tags, and delete
- View connection information such as:
 - The connection's OCID
 - The compartment in which the connection was created
 - When the connection was created and updated
 - Encryption key management

- Network information
 - * View and edit the connection's subnet

 **Note:**

If you change the subnet, the connection's Ingress IPs and Private IP also changes.

- * Traffic routing method, either Shared or Dedicated
 - * The **Ingress IP** show one or more IP addresses. The connection endpoint for OCI GoldenGate originates from one of these IP addresses. You must ensure that the appropriate subnet security rules are in place to allow connectivity from these IP addresses to the database's private IP.
 - * The **Private IP** is the private connection endpoint.
 - * For MySQL databases, OCI GoldenGate generates the Host value depending on the IP you provide during creation. The Private IP gets rewritten in the format, `ip-10-0-0-0.ocigsvc.oracle.vcn.com`.
- Connection type and connectivity information specific to that type
 - View, assign, unassign, and test connections
 - View, add, edit, or remove network security groups (NSGs) for connections with Dedicated endpoints

 **Note:**

Connections with Shared endpoints inherit NSG settings from the assigned deployment. Adding network security groups (NSGs) gives you fine grained control over traffic between a deployment and the source or target systems within your subnet. [Learn more](#).

- View the connection's work requests

Edit a connection

 **Note:**

For Autonomous Database connections, editing and saving the connection automatically retrieves the latest wallet version.

 **WARNING:**

You must only create and edit connections in the Oracle Cloud console. Refrain from creating or editing connections in the Credentials screen of the deployment console. Updates are automatically synced to the deployment from the Oracle Cloud console.

To edit a connection:

1. From the Connections page, select a connection to edit.

You can also select **Edit** from the connection's Action (ellipsis icon) menu, and then skip to step 3.

2. On the Connection Details page, click **Edit**.
3. In the Edit Connections panel, edit the values you want to update.

 **Note:**

You can update the Traffic routing method and subnet selection, if needed. If you change the subnet, the connection's Ingress IPs and Private IP also change.

4. Click **Save changes**.

The changes are automatically synchronized to the OCI GoldenGate deployment that it's assigned to.

Assign a deployment to a connection

When you're on the Connections details page, you can assign deployments to the connection.

To assign a deployment to a connection:

1. On the Connections details page, under **Resources**, click **Assigned deployments**.
2. Click **Assign deployment**.
3. In the Assign deployments dialog, select a connection from the dropdown. Click **Change Compartment**, to select a connection from a different compartment.
4. Click **Assign deployment**.

The selected connection appears in the Assigned deployments list. You can also view and manage this relationship from the Deployment details page, under **Assigned connections**.

Unassign a deployment

You can remove, or unassign, a connection from a deployment. First ensure that there are no active processes running in your deployment that involve the connection you want to remove.

To unassign a deployment:

1. On the Connection details page, under **Resources**, click **Assigned deployments**.
2. From the Actions (ellipsis icon) menu for the deployment you want to unassign, select **Unassign**.
3. In the Unassign deployment dialog, confirm that you want to unassign the connection from this deployment, and then click **Unassign deployment**.

The deployment no longer appears in the Assigned deployments list. Unassigning a connection from a deployment doesn't delete the connection. If you want to delete the connection see, Delete a connection.

Test connections

After you assign connections to a deployment, test the connectivity to ensure the deployment can connect to them.

Before you test a connection, ensure that you first create and assign a connection to the deployment.

To test an assigned connection:

1. On the connection details page, under **Resources**, click **Assigned Deployments**.

You can also test the connection from the Assigned connections list on a deployment's details page.

2. In the list of Assigned deployments, from the Action (three dots) menu for the connection you want to test, select **Test connection**.

The Test connection dialog opens and displays a confirmation message, if successful, and an error message, if unsuccessful. If an error message appears, then return to the connection and your settings.

3. Click **Close**.

Move a connection

To move a connection:

1. On the Connection Details page, click **Move resource**. You can also select **Move connection** from the connection's Actions (ellipsis icon) menu on the Connections page.
2. In the Move to another compartment dialog, select the compartment in which to move the connection.
3. Click **Move resource**.

Manage tags for a connection

Tags help you locate resources within your tenancy. You can add and view a connection's tags from the Connections page and from the connection details page.

On the Connections page, from a connection's Action menu (ellipsis icon), select **Add Tags** or **View Tags**.

On the connections details page, you can select **Add Tags** from the **More Actions** menu, or click the **Tags** tab to view and edit tags.

Delete a connection

Before you delete the connection, ensure that there are no processes running in the deployment that involve the connection. Consider unassigning the connection from deployments before you delete it.

To delete a connection:

1. On the Connections page, select the connection you want to delete. You can also select **Delete** from the connection's Actions (ellipsis) menu in the Connections list, and then skip to Step 3.
2. On the Connections details page, click **Delete**.
3. In the Delete connection dialog, confirm that you want to delete the connection, and then click **Delete**.

Known issues

Action Required for Autonomous Databases that Use mTLS Authentication

When an Autonomous Database wallet is rotated, the OCI GoldenGate connection to this database must be refreshed to retrieve the latest wallet information.

For more information see, [My Oracle Support \(MOS\) Document 2911553.1](#).

To refresh an Autonomous Database connection: Edit and save the connection to the Autonomous Database (Autonomous Transaction Processing or Autonomous Datawarehouse). Saving the connection automatically downloads and refreshes the wallet. No other changes to the connection is needed.

To verify:

1. Launch the deployment console for a deployment that uses the Autonomous Database connection.
2. In the deployment console, open the navigation menu, and then click **Configuration**.
3. On the Credentials screen, observe the Autonomous Database connection string. Before the wallet is refreshed, the connection string looks like the following:

```
ggadmin@(DESCRIPTION=(TRANSPORT_CONNECT_TIMEOUT=3) (CONNECT_TIMEOUT=60)
(RECV_TIMEOUT=120) (retry_count=20) (retry_delay=3) (address=(protocol=tcps)
(port=1522) (host=adb.us-phoenix-1.oraclecloud.com) )
(CONNECT_DATA=(COLOCATION_TAG=ogginstance) (FAILOVER_MODE=(TYPE=SESSION)
(METHOD=BASIC) (OVERRIDE=TRUE)) (service_name=<adb-
servicename>_low.adb.oraclecloud.com) )
(security=(MY_WALLET_DIRECTORY="/u02/connections/
ocid1.goldengateconnection.oc1.phx.<ocid>/wallet")
(SSL_SERVER_DN_MATCH=TRUE) (ssl_server_cert_dn="CN=adwc.uscom-
east-1.oraclecloud.com,
OU=Oracle BMCS US, O=Oracle Corporation, L=Redwood City,
ST=California,
C=US")))
```

After the wallet is refreshed, the connection string is updated to look like the following:

```
ggadmin@(DESCRIPTION=(TRANSPORT_CONNECT_TIMEOUT=3) (CONNECT_TIMEOUT=60)
(RECV_TIMEOUT=120) (retry_count=20) (retry_delay=3) (address=(protocol=tcps)
(port=1522) (host=adb.us-phoenix-1.oraclecloud.com) )
(CONNECT_DATA=(COLOCATION_TAG=ogginstance) (FAILOVER_MODE=(TYPE=SESSION)
(METHOD=BASIC) (OVERRIDE=TRUE)) (service_name=<adb-
servicename>_low.adb.oraclecloud.com) )
(security=(MY_WALLET_DIRECTORY="/u02/connections/
```

```
ocid1.goldengateconnection.oc1.phx.<ocid>/wallet")  
(SSL_SERVER_DN_MATCH=TRUE)(ssl_server_dn_match=yes)))
```

Managing deployments

Learn to edit, scale, stop, start, and delete a deployment in the Oracle Cloud console.

View deployment details

Select a deployment from the Deployments page to view its details. On the deployment details page, you can:

- View the deployment's status, which can be one of the following:
 - Creating
 - Updating
 - Active
 - Inactive
 - Deleting
 - Deleted
 - Failed
 - Needs Attention

 **Note:**

If the deployment's status is Needs Attention, then the deployment health is less than 100%. Learn about OCPU management and billing.

- Perform actions on the deployment, such as:
 - Edit
 - Stop or Start
 - Scale
 - Collect diagnostics for Data replication deployments
 - Upgrade
 - Move
 - Add tags
 - Delete
- Launch the deployment console.
- Launch Admin client for Data replication deployments.
- View deployment information, such as the deployment's OCID, when it was created and last updated
- Set up contextual notifications and stay informed of deployment events

- View the Data replication deployment's storage utilization. Storage utilization limit indicates how much file system space the deployment is currently using.

 **Note:**

You have starting soft limit of 250 GB of space per deployment OCPU, and a hard limit of 500 GB per base OCPU. For example, for 4 base OCPUs, the soft limit is 1 TB and the hard limit 2 TB. If you exceed the hard limit, the service may be limited in functionality and performance.

Create alarms to notify you of this event. Manage trail files to free up space.

- View and schedule the deployment's maintenance window.
- View GoldenGate instance information, such as:
 - Deployment type
 - Show/Copy its console URL
 - Build version
 - Edit a deployment username and/or password.
- View Network information, such as:
 - View/Edit the deployment's subnet

 **Note:**

If you change the subnet, the deployment's Ingress IPs and Private IP also change.

- Public IP address
- Private IP address
- Ingress IPs, if assigned connections' Traffic routing method is Shared.
- Load balancer information, if you enabled GoldenGate console public access.

 **WARNING:**

While you can manage this resource, ensure that you don't delete the load balancer while your deployment is still in use.

- View deployment metrics.
- Add and view Network Security Groups.

 **Note:**

Adding network security groups (NSGs) gives you fine grained control over where the deployment can be accessed from within your subnet. [Learn more.](#)

- Create and view deployment backups.
- View upgrade history.
- View, assign, unassign connections, and test connections.
- View trail files details such as size, sequence, producers, and consumers.

 **Note:**

In Data replication deployments, Trail files can build up over time and are major contributors to the Storage utilization number you see under Deployment information. Use this information to manage trail files.

- Import and export Master encryption key wallets for Data replication deployments.
- View the status of work requests and any log messages, error messages, and resources associated with them. Deployment operations that create work requests include:
 - Create
 - Update
 - Delete
 - Move
 - Restore
 - Stop
 - Start
 - Patch
- Enable/disable process and error logs for Data replication deployments.

 **Note:**

You can also enable, disable, and view Logs from the OCI Logging service in the Oracle Cloud console.

Edit a deployment

To edit a deployment:

1. On the Deployments page, select a deployment, and then click **Edit** from the deployment details page.

You can also select **Edit** from the Actions (three dots) menu of the deployment you want to edit.
2. In the Edit Deployment dialog, you can update the following fields:
 - Name
 - Description
 - Enable GoldenGate console public access

- FQDN Prefix
 - Certificate and private key pair (keep, remove, or replace)
3. Click **Save Changes**.

Edit a deployment username

You can edit the credential store used to log into the deployment console.

If you have yet to change the credential store to Oracle Cloud Infrastructure Identity and Access Management (OCI IAM), ensure that you first set up your Vault. Learn more about [Vault](#) and [managing secrets](#).

1. On the deployment details page, in the GoldenGate section of the Deployment information card, click **Edit** next to **Username**.

The Edit username panel opens.

2. In the Edit username panel, you can:
 - Change the credential store from GoldenGate to Oracle Cloud Infrastructure Identity and Access Management (OCI IAM).
 - Change the username, if your credential store is GoldenGate.
 - Select a different password secret, if your credential store is GoldenGate.

Note:

To edit the contents of the existing password secret, see [Edit a password secret](#).

- Create a new password secret. To create a new password secret:
 - a. Click **Create password secret**.
 - b. In the Create secret panel, enter a name for the secret, and optionally, a description.
 - c. Select a compartment from the **Compartment** dropdown in which to save your secret.
 - d. Select a vault in the current compartment, or click **Change compartment** to select a vault in a different compartment.
 - e. Select an **Encryption key**.

Note:

Only AES keys, Software protected keys, and HSM keys are supported. RSA and ECDSA keys are not supported for GoldenGate password secret keys.

- f. Enter a password 8 to 30 characters in length, containing at least 1 uppercase, 1 lowercase, 1 numeric and 1 special character. The special characters must not be '\$', '^' or '?'.
 - g. Confirm the password.


```
allow group ContextualNotificationsUsers to manage ons-topics in tenancy
allow group ContextualNotificationsUsers to manage cloudevents-rules in
tenancy
```

To set up contextual notifications for a deployment:

1. On the Deployments page, select the deployment you want to set up notifications for.
2. On the deployment details page, switch to the **Notifications** tab.

When no notifications exist for the deployment, the Notifications tab lists available quick start templates. You can select one of the templates, or click **Create notification**.

3. In the Create Notification panel, in the **Topics and subscriptions** section, create or select an existing topic.
4. Select at least one subscription protocol:
 - Email, and then enter a valid email address.
 - Slack, and then enter the Slack endpoint.
Sends a message to the specified Slack channel by default, when you publish a message to the subscription's parent topic.

Endpoint format (URL):

```
https://hooks.slack.com/services/<webhook-token>
```

The `<webhook-token>` should contain two forward slashes (/). Query parameters aren't permitted in URLs. To create an endpoint for a Slack subscription (using a webhook for the Slack channel), see the [Slack documentation](#).

- SMS, and then select the country code and enter the phone number.
5. (Optional) Modify the default settings (event type or alarm severity) listed above **Topic and Subscriptions**.
 6. Click **Create notification**.
 7. Confirm the new subscriptions, if needed. For more information, see [Confirming a subscription](#).

Messages are sent to the contact information entries whenever the condition of the notification is satisfied.

Scale a deployment

You can scale a deployment up or down depending on the amount of Oracle Compute Units (OCPU) you need.

When you enable auto scaling, the deployment can scale up to three times more memory than the number of OCPUs currently shown in the Scale dialog. One OCPU is equivalent to 16 GB memory. If your workload requires additional OCPUS, then the deployment automatically uses the resources without any manual intervention required when auto scaling is enabled.

To see OCPU usage, you can view the OCPU Consumption graph in the Metrics section of the Deployment Details page in the Console.

To scale a deployment:

1. On the Deployments page, from the Actions (three dots) menu of the deployment you want to scale, select **Scale**.

You can also click **Scale** on the deployment's Details page.

2. For OCPU Count, enter the number of OCPUs between 1 and 24.
3. (Optional) Enable **Auto Scaling**.

Enabling auto scaling allows the service to scale up to three times the designated number of OCPUs, for a maximum of 24 OCPUs.

4. Click **Save Changes**.

Your deployment restarts to reflect the changes made.

Collect diagnostics

Collect diagnostics to analyze or share information about your OCI GoldenGate deployment. The information collected can be shared with My Oracle Support if you encounter any issues.



Note:

This feature applies only to Data replication deployments.

Before you collect diagnostics, ensure that you [create an OCI Object Storage bucket](#).

To collect deployment diagnostics:

1. On the Deployments page, select the deployment for which to collect diagnostics.
2. On the deployment's Details page, from the More actions menu, select **Collect diagnostics**.
3. In the Collect diagnostics panel, complete the following fields, and then click **Collect diagnostics**:
 - a. From the **Bucket** dropdown, select the bucket in which to save the diagnostics file. If you want to select a bucket in a different compartment, click **Change Compartment**.
 - b. For **Diagnostics name prefix**, enter a short name or a few characters to prefix the diagnostics file name.
 - c. (Optional) Select the Start date from which to collect system logs.
 - d. (Optional) Select the End date to which to collect system logs.



Note:

Diagnostics collected for the OCI GoldenGate deployment contains GoldenGate logs for the entire lifespan of the deployment and are independent of the start and end dates selected.

After you click Collect diagnostics, a new field named **Diagnostics** displays under the GoldenGate section of the Deployment information card. It can take several minutes for the diagnostics zip file to become available to download. When it is available, a **Download** link appears.

Stop a deployment

When you stop a deployment, Oracle Cloud Infrastructure GoldenGate also stops. You won't be able to access the OCI GoldenGate deployment console while the deployment is stopped and you won't be billed until you restart it.

To stop a deployment:

1. From the Deployments page, select a deployment, and then click **Stop** on the deployment details page.

You can also select **Stop** from the Actions (three dots) menu of the deployment you want to stop on the Deployments page.

2. In the Stop Deployment dialog, click **Stop**.

When you stop a deployment, Oracle Cloud Infrastructure GoldenGate stops all active tasks. You can restart the deployment from the Deployments or deployment details pages.

Start a deployment

After the deployment is created, Oracle Cloud Infrastructure GoldenGate is automatically started. If you stop a deployment, you can restart it using the Start option in the deployment's Actions (three dots) menu. When you start the deployment, billing also resumes.

To start a deployment:

1. From the Deployments page, select a deployment, and then click **Start** on the deployment details page.

You can also select **Start** from the Actions (three dots) menu of the deployment you want to start on the Deployments page.

2. In the Start Deployment dialog, click **Start**.

The deployment starts and you can now launch the Deployment Console. Oracle resumes billing for the amount of Oracle compute units (OCPU) used.

You can also configure Extracts and Replicats to automatically start when you start the deployment. For more information, see [Configure managed processes](#).

Move a Deployment

You can move a deployment from one compartment to another compartment.

To move a deployment:

1. From the Deployments page, select a deployment, and then in the **Actions** menu, select **Move Resource** from the deployment details page.

You can also select **Move Resource** from the Actions (three dots) menu of the deployment you want to move on the Deployments page.

2. In the Move Resource to a Different Compartment dialog, select the compartment to move the deployment to from the dropdown.
3. Click **Move Resource**.

After you move the deployment to the new compartment, inherent policies apply immediately and may affect access to the deployment through the Console. For more information, see [Managing Compartments](#).

Manage tags for a deployment

Tags help you locate resources within your tenancy. You can add and view a deployment's tags from the Deployments page and from the deployment details page.

On the Deployments page, from a deployment's Actions (three dots) menu, select **Add Tags** or **View Tags**.

On the deployment details page, you can select **Add Tags** from the **More Actions** menu, or click the **Tags** tab to view and edit tags.

[Learn more about tagging.](#)

Delete a deployment

When you delete a deployment, all active Oracle GoldenGate tasks within that deployment stop. Deleting a deployment doesn't remove references, such as checkpoint tables and Extract information, from the databases the deployment used. After a deployment is deleted, it *cannot* be restored.

To delete a deployment:

1. On the Deployments page, select a deployment, and then from the **More Actions** menu, select **Delete**.

You can also select **Delete** from the Actions (three dots) menu for the deployment you want to delete on the Deployments page.

2. In the Delete Deployment dialog, click **Delete**.

You may need to manually remove files that remain in the source and target databases after a deployment is deleted. For more information, see [Files to be removed manually](#).

Assign a connection to a deployment

Ensure that you have connections created for your source and target technologies.

To assign a connection to a deployment:

1. On the deployment details page, under **Resources**, click **Assigned connections**.
2. Click **Assign connection**.
3. In the Assign connection dialog, select a connection from the dropdown. If you want to select a connection from a different compartment, click **Change Compartment**.
4. Click **Assign connection**.

The selected connection appears in the Assigned connections list. You can also view and manage this relationship from the Connection details page under **Assigned deployments**.

Test connections

After you assign connections to a deployment, test the connectivity to ensure the deployment can connect to them.

Before you test a connection, ensure that you first create and assign a connection to the deployment.

To test an assigned connection:

1. On the deployment details page, under **Resources**, click **Assigned Connections**.

You can also test the connection from the Assigned deployments list on a connection's details page.

2. In the list of Assigned connections, from the Action (three dots) menu for the connection you want to test, select **Test connection**.

The Test connection dialog opens and displays a confirmation message, if successful, and an error message, if unsuccessful. If an error message appears, then return to the connection and your settings.

3. Click **Close**.

Unassign a connection

You can remove, or unassign, a connection from a deployment. First ensure that there are no active processes running in your deployment that involve the connection you want to remove.

To unassign a connection:

1. On the Deployment details page, under **Resources**, click **Assigned connections**.
2. From the Actions (ellipsis icon) menu for the connection you want to unassign, select **Unassign**.
3. In the Unassign connection dialog, confirm that you want to unassign the connection from this deployment, and then click **Unassign connection**.

The connection no longer appears in the Assigned connections list. Unassigning a connection does not delete the connection. If you want to delete the connection, see Delete a connection.

Manage deployment users

Managing GoldenGate users depends on whether you're using an Identity Access Management (IAM) enabled tenancy or not. In IAM enabled tenancies, you use your Oracle Cloud account to access the deployment console. In non-IAM enabled tenancies, each deployment has its own set of users.

 **Note:**

This information applies only to Data replication deployments.

In IAM-enabled tenancies

Identity Access Management (IAM) enabled tenancies let you create identity domains to manage users and roles, federate and provision users, secure application integration through Oracle Single Sign-On (SSO) configuration, and SAML/OAuth based Identity Provider administration.

 **Note:**

These features can only be used in tenancies already migrated to Identity Access Management (IAM) with Identity domains. Once you select IAM as your deployment's credential store, you won't be able to change back.

Configure Identity domains for OCI GoldenGate

You can create identity domains in IAM-enabled tenancies. The following steps describe how to create groups of users and configure password policies for your domain.

The domain settings mentioned here are specific to OCI GoldenGate. [Learn more about identity domains and how to create one.](#)

1. In the Oracle Cloud console navigation menu, select **Identity & Security**, and then under **Identity**, click **Domains**.
2. From the list of Domains, select your identity domain.
3. On your identity domain overview page, from the Identity domain menu, click **Groups**.
4. Create the following groups to map to GoldenGate roles:
 - GGS_Administrator
 - GGS_Security
 - GGS_Operator
 - GGS_User

 **Note:**

GoldenGate roles are as follows:

- **Administrator:** Grants full access to the user, including the ability to alter general, non-security related operational parameters and profiles of the OCI GoldenGate deployment service.
- **Security:** Grants administration of security related objects and invoke security related service requests. This role has full privileges.
- **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 OCI GoldenGate deployment services.
- **User:** Allows information-only service requests, which do not alter or effect the operation of either the OCI GoldenGate deployment services.

5. Select the users to add to the group, and then click **Create**.

 **Note:**

Each group must be assigned at least one user. [Learn more about groups.](#)

6. Set the **Access signing certificate** option.
 - a. From the Identity domain menu, click **Settings**.
 - b. For Access signing certificate, select **Configure client access** to allow clients to access the tenant signing certificate and the SAML metadata without logging in to the identity domain.
 - c. Click **Save changes**.
7. Specify the password policy for your Identity domain:
 - a. From the Domain settings menu, click **Password policy**.
 - b. On the Password policies page, you can edit the default password policy or add a new one.

In non-IAM enabled tenancies

In non-IAM enabled tenancies, deployment user management occurs within the OCI GoldenGate deployment console. Each OCI GoldenGate deployment can have its own set of users.

Add a user to a deployment

To add a user:

1. Launch the OCI GoldenGate deployment console from the deployment details page.
2. Log in to the OCI GoldenGate deployment console as the Oracle GoldenGate Administrator user.

 **Note:**

The Administrator user was created when the deployment was created.

3. Open the OCI GoldenGate deployment console navigation menu, and then click **Administrator**.
4. Click **Add User** (plus icon).
5. For **Name**, enter a unique user name.

 **Note:**

The user name must start with an alphabetic character and contain only alphanumeric characters. Symbols that can be used are: at sign (@), period (.) , dash(-), comma(,), underscore(_), number sign(#), dollar sign(\$), plus sign (+), backslash (\), slash (/), equal sign (=), less than sign (<), or greater than sign(>)

6. For **Role**, select one of the following roles:
 - **User**: Allows information-only service requests, which do not alter or effect the operation of either the OCI GoldenGate deployment services.
 - **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 OCI GoldenGate deployment services.
 - **Administrator**: Grants full access to the user, including the ability to alter general, non-security related operational parameters and profiles of the OCI GoldenGate deployment service.
 - **Security**: Grants administration of security related objects and invoke security related service requests. This role has full privileges.
7. (Optional) For **Description**, enter a short description.
8. For **Type**, select **Basic** from the dropdown.

 **Note:**

Certificate type user accounts is not currently supported in OCI GoldenGate.

9. Enter a password, and then enter it again to verify.

 **Note:**

The password must be 8 to 30 characters and contain at least 1 uppercase, 1 lowercase, 1 numeric and 1 special character. The special characters must not be '\$', '^' or '?'.

10. Click **Submit**.

The deployment user account appears in the Users list. You can edit or delete the user from the Actions column.

Edit a deployment user

When you edit a deployment user, you can only change the Info and Password values. Certificate type user accounts are not currently supported by OCI GoldenGate.

To edit a user:

1. Launch the OCI GoldenGate deployment console from the deployment details page.
2. Log in to the OCI GoldenGate deployment console as the Oracle GoldenGate Administrator user.

 **Note:**

The Administrator user was created when the deployment was created.

3. Open the OCI GoldenGate deployment console navigation menu, and then click **Administrator**.
4. For the user account you want to edit, click **Edit user**.
5. Make your changes, and then click **Submit**.

 **Note:**

Passwords must be 8 to 30 characters and contain at least 1 uppercase, 1 lowercase, 1 numeric and 1 special character. The special characters must not be '\$', '^' or '?'.

If you changed the user account password, ensure that you also update the user credentials for any Oracle GoldenGate processes that involve this user.

Delete a deployment user

Oracle recommends that you periodically review deployment user accounts and remove inactive accounts.

To delete a user from a deployment:

1. Launch the OCI GoldenGate deployment console from the deployment details page.
2. Log in to the OCI GoldenGate deployment console as the Oracle GoldenGate Administrator user.

 **Note:**

The Administrator user was created when the deployment was created.

3. Open the OCI GoldenGate deployment console navigation menu, and then click **Administrator**.

4. In the Users list, locate the user to delete, and then click **Delete user** (trash icon) in the Action column associated with that user.
5. In the Confirm Deletion dialog, verify that this is the user correct user you want to delete, and then click **OK**.

The user is removed from the Users list.

Managing Trail files

OCI GoldenGate Trail files quickly add up over time. Without purge tasks in place to manage these trail files, daily backups will take exponentially longer to complete and use compute resources that could otherwise be used elsewhere.



Note:

This information applies only to Data replication deployments.

View Trail files

You can view Trail files details such as size, sequence, producers, and consumers on the deployment details page.

To view Trail files:

1. From the GoldenGate Overview page, click **Deployments**.
2. On the Deployments page, select a deployment to view its details, or select **View details** from the deployment's Action menu (ellipsis icon).
3. On the Deployment details page, under **Resources**, click **View Trail files**.

Use this information to identify Trail files that are no longer used or needed, and then create Purge tasks to manage them.

Purge Trail files

OCI GoldenGate Trail files quickly add up over time. Without purge tasks in place to manage these Trail files, daily backups will take exponentially longer to complete and use compute resources that could otherwise be used elsewhere.

Oracle recommends that you create a backup first, before purging Trail files. See [Create a manual backup](#). You can then download the backup and review its contents (`<deployment-name>/var/lib/data/`).

To purge OCI GoldenGate Trail files:

1. In the OCI GoldenGate deployment console, review the following processes that generate Trail files, and take note of the Trail files you want to keep:
 - In the Administration Service, review the details of each Extract and Replicat. (Click the process name, and then select **Details**.)
 - In the Distribution Service, review the details of each Distribution Path (if any).
 - In the Receiver Service, review the details of each Receiver Path (if any).

All other Trail files not on your list can be purged.

2. Set up Purge Tasks to clean up unused Trail files in the deployment console.
 - a. In the OCI GoldenGate deployment console, open the navigation menu (hamburger icon) for the Administration Service, and then select **Configuration**.
 - b. On the **Configuration** page, click **Tasks**, and then click **Purge Trails**.
 - c. Click **Add Purge Trails Task** (plus icon).
 - d. Under Create a new Purge Trails task, complete the following fields, and then click **Submit**:
 - i. For **Operation Name**, enter a name for the purge task.
 - ii. For **Trail**, enter a name of a Trail file, and then click **Add Trail** (plus icon). Repeat this step to add more Trail files. **Selected Trails** populates with the names of the Trail files as you add them.
 - iii. Keep **Use Checkpoints** enabled if you want to purge after all Extract and Replicat processes are finished with the file(s), as indicated by checkpoints. Disabling this option allows purging without considering checkpoints and a minimum of one file (if no MIN value is specified) or the number of files specified with MIN are kept.

 **Note:**

Orphan Checkpoint files can't be deleted in the deployment console. Use REST APIs to delete orphan Checkpoint files.

- iv. For **Keep Rule**, specify the Hours, Days, or Number of Files to keep.
- v. For **Purge Frequency**, specify the frequency to run this purge task.

You can add more Purge Tasks or disable them as needed.

3. Clean up unused Trail files using Admin Client.
 - a. On the deployment details page, click **Launch Admin Client**.
If connecting to an OCI GoldenGate deployment with a private endpoint, follow the instructions in Connect to Admin Client through Cloud Shell.
 - b. Run the following command to purge trail files no longer used by Extracts:

```
purge exttrail <trail-file-name>
```

For more information, see PURGE EXTTRAIL in the *GoldenGate Command Line Reference* guide.

4. Clean up unused Trail files using REST APIs.
 - a. Open Cloud Shell.
 - b. In Cloud Shell, run the following REST API call to the OCI GoldenGate deployment. Ensure that you replace the placeholders with your deployment's actual values.

```
curl <console url> -u <username>:<password> -X POST -H 'Content-Type: application/json' -d '{ "trails": [ { "name": "<one or two character trail prefix>" }, {"keep": [ { "type": "min", "units": "<days | hours | files>", "value": <number of units" } ], "name": "purge", "purgeType": "trails", "useCheckpoints": true | false } ] }
```

To ensure all Trail files are deleted, you can do one of the following:

- On the deployment details page in the Oracle Cloud console, under **Resources**, click **Trail files** and then click **Refresh** to review the list of Trail files.
- Create another manual backup and check the contents of the backup for the Trail files you wanted to keep. Note the aggregate size reduction and speed that the backup now completes.

Manage master encryption key wallets

Use master encryption keys to encrypt trail files distributed to other GoldenGate deployments. You can then import and export master encryption key wallets to use with other source and target OCI GoldenGate deployments.



Note:

This information applies only to Data replication deployments.

If a master key is created in Oracle GoldenGate, then each time GoldenGate creates a trail file, it automatically generates a new encryption key that encrypts the trail contents. The master key encrypts the encryption key.

Before you begin

Ensure that you have the following:

- Access to the [Vault service](#) and a [Vault created](#)



Note:

A virtual private vault is not required.

- Added the minimum required policies to for OCI GoldenGate to use the Vault service
- [A master encryption key created in your Vault](#)



Note:

Only AES, software protected keys, or HSM keys are supported. RSA and ECDSA are not supported.

Add a master key in the deployment console

To add a master key in the GoldenGate deployment console:

1. Launch the GoldenGate deployment console from the deployment details page.
2. Log in as the GoldenGate admin user.

3. After you log in, open the navigation menu, click **Configuration**, and then click **Key Management**.
4. On the Key Management page, for Master Keys, click **Add Master key** (plus icon).
A new master key appears in the list.

Export a master encryption key wallet from an OCI GoldenGate deployment

If a master key is added in the source deployment, ensure that you export it and import it into the target deployment.

To export a master encryption key wallet:

1. On the Deployments page, select the deployment from which to export the master encryption key wallet.
2. On the deployment details page, under Resources, click **Master encryption key actions**.
3. Click **Export**.
4. In the **Export** dialog:
 - a. For Name, enter a name for the master encryption key wallet.
 - b. (Optional) Enter a description to help distinguish it from others in the wallet list.
 - c. For **Vault in <compartment-name>**, select the vault in which to export the master encryption key wallet. Click **Change compartment** to select a different compartment.
 - d. For **Encryption key in <compartment name>**, select the appropriate encryption key to use. Click **Change compartment** to select a different compartment.
5. Click **Export**.

Export a master key encryption wallet from an on premise Oracle GoldenGate instance

If a master key is added to a source (on premise or Marketplace) Oracle GoldenGate instance, ensure that you base64 encode the `cwallet.sso` and then copy it into an OCI Vault secret.

To export a master encryption key wallet from an on premise Oracle GoldenGate instance:

1. SSH into your on premise Oracle GoldenGate instance.
2. Change directories to the location in which the wallet (`cwallet.sso`) resides.

 **Note:**

Oracle recommends making a copy of `cwallet.sso` to work with.

3. Base64 encode the `cwallet.sso` using the following command:

```
base64 -w 0 cwallet.sso
```

4. Copy the output string.
5. In the **Oracle Cloud console**, open the navigation menu, select **Identity & Security**, and then select **Vault**.
6. On the Vaults page, select your vault.
7. On the Vault details page, under **Resources**, click **Secrets**, and then click **Create Secret**.
8. In the Create Secret panel, complete the fields as follows:
 - a. For **Create in Compartment**, select the compartment in which to create the Secret.
 - b. Enter a **Name** for the secret.
 - c. (Optional) Enter a **Description** for the secret.
 - d. For **Encryption Key in <compartment-name>**, select the master encryption key created in the *Before you begin* steps. Click **Change compartment** to select a master encryption key located in a different compartment.
 - e. For **Secret type template**, select **Plain-Text**.
 - f. For **Secret contents**, paste the cwallet.sso base64 encoded string from step 3.
9. Click **Create Secret**.

The Secret appears in the Secrets list. You can now import the master encryption key wallet to the target OCI GoldenGate deployment, and select this Secret.

Import a master encryption key wallet to a deployment

To import a master encryption key wallet:

1. On the Deployments page, select the deployment in which to import the master encryption key wallet.
2. On the deployment details page, under Resources, click **Master encryption key wallet actions**.
3. Click **Import**.
4. In the **Import** dialog:
 - a. For **Wallet secret in <compartment-name>**, select the wallet secret to import. Click **Change compartments** to select a wallet secret from a different compartment.
 - b. (Optional) Select **Backup existing wallet to ...**
If selected, then under **Backup wallet**:
 - i. For **Name**, enter a name for the backup wallet.
 - ii. (Optional) Enter a description.
 - iii. For **Encryption key in <compartment-name>**, select the encryption key to use. Click **Change compartment** to select an encryption key in a different compartment.
5. Click **Import**.

Import a master encryption key wallet to an on premise GoldenGate instance

Ensure that you exported the source OCI GoldenGate deployment's master encryption key wallet.

To import a master encryption key wallet to an on premise GoldenGate instance:

1. On the OCI GoldenGate Deployments page, select the source deployment.
2. On the deployment details page, under **Resources**, click **Master encryption key wallet actions**.
3. In the Master encryption key wallet actions list, select the exported wallet. You're brought to the Secret details page in your Vault.
4. On the Secret details page, under **Versions**, open the **Action menu** (ellipsis icon), and then select **View Secret contents**.
5. In the View Secret contents dialog, select **Show decoded Base64 digit**.
6. Copy the contents of the text area.
7. SSH into your on premise or Marketplace Oracle GoldenGate instance.
8. Create a new text file (`vi` or other text editor) and then paste the Secret contents into the file.
9. Run the Base64 command on the file you created (ensure that you replace `<filename>` with the name of your text file):

```
base64 -d <filename> > cwallet.sso
```

10. Copy or move `cwallet.sso` to the GoldenGate wallet directory.

You can now add and run a Replicat to receive the encrypted Trail file sent from the source OCI GoldenGate deployment.

Managing Truststore certificates

Learn to add and delete a Truststore certificate in the Oracle Cloud console.



Note:

This information applies only to Data replication deployments.

What is a Truststore certificate?

GoldenGate deployments can communicate with each other in a client-server manner. Typically, the deployment initiating communication is the client, and the deployment receiving communication is the server.

A Truststore certificate is a client-side asset that functions as a repository of certificates from trusted Certificate authorities (CA). It is used to verify the identity of the server's certificate being presented to the client.

Add a Truststore certificate

Select a deployment from the Deployment list to view its details. On the Deployment details page, you can:

1. On the Deployment details page, under Resources, click **Truststore certificate**.
2. Click **Add certificate**.
3. In the Add Truststore certificate dialog:
 - a. For **Key**, enter an identifier key for the Truststore certificate.

 **Note:**

This identifier key must be unique within the scope of the deployment. It must also be 1 to 32 alphanumeric characters long, starting with a letter.

- b. For **Truststore certificate**, upload a certificate by dragging the file into the location or click **select one** and choose the certificate you want to upload.
4. Click **Add**.

Delete a Truststore certificate

When you're on the Deployment details page, you can delete a Truststore certificate.

 **WARNING:**

If no other certificate exists for the server deployment, subsequent connections to the server deployment will fail.

To delete a Truststore certificate:

1. On the Deployment details page, under Resources, click **Truststore certificate**.
2. In the list of Truststore certificates, from the Action (three dots) menu for the connection you want to remove, select **Delete**.
3. In the Delete certificate dialog, click **Delete**.

Manage deployment backups

Deployment Backups can help you troubleshoot issues, clone deployments, and restore deployments. Deployment Backups are automatically created daily. You can also create manual backups at any time, and save them to your Oracle Object Storage tenancy.

Creating a Manual Backup

To create a backup:

1. On the deployment details page, under **Resources**, click **Deployment backups**.
2. Click **Create backup**.
3. In the Create backup panel, enter a name for the backup.
4. From the **Compartment** dropdown, select the compartment of your tenancy's Oracle Object Storage.
5. From the **Oracle Object Storage in <compartment_name>** dropdown, select the Object Storage bucket where your backup will be stored.
6. For **Object Name**, enter a name for the backup as it'll appear in Object Storage. Ensure to add the file extension `.xz` at the end of the object name.

 **Note:**

You can then download the manual backup and use a client that supports `.xz` file types to unzip it.

7. Click **Create backup**.

Manual backup directory structure

You can create a manual backup of your OCI GoldenGate and save it to your Oracle Object Storage tenancy. From there, you can download the manual backup locally, and then access a full directory of the GoldenGate deployment's files, including log and trail files.

The directory structure of a manual backup appears similar to the following:

- `bin`
- `cfgtoollogs`
- `deinstall`
- `diagnostics`
- `include`
- `install`
- `inventory`
- `jdk`
- `jlib`
- `lib`
 - `instantclient`
 - `sql`
 - `utl`

- OPatch
- oraInst.loc
- oui
- srvnm

The following table describes the key Oracle GoldenGate directories and the variables used to reference them. When you see these variables in an example or procedure, replace the variable with the full path to the corresponding directory path in your manual backup directory.

Directory Name	Variable	Description	Default Directory Path
Oracle GoldenGate home	OGG_HOME	The Oracle GoldenGate home that is created on a host computer is the directory that you choose to install the product. This read-only directory contains binary, executable, and library files for the product.	/ogg_install_location
Deployment configuration home	OGG_CONF_HOME	The location where each deployment information and configuration artifacts are stored.	/ogg_deployment_location/etc/conf
Deployment security home	OGG_SSL_HOME	The location where each deployment security artifacts (certificates, wallets) are stored.	/ogg_deployment_location/etc/ssl
Deployment data home	OGG_DATA_HOME	The location where each deployment data artifacts (trail files) are stored.	/ogg_deployment_location/var/lib/data
Deployment variable home	OGG_VAR_HOME	The location where each deployment logging and reporting processing artifacts are stored.	/ogg_deployment_location/var
Deployment etc home	OGG_ETC_HOME	The location where your deployment configuration files are stored including parameter files.	/ogg_deployment_location/etc

Viewing deployment backup details

From the Deployment Backups list, click a backup to view its details.

On the Deployment Backup Details page, you can:

- Restore
- Create Clone
- Move it to another compartment
- Add Tags
- Delete it
- View Deployment Backup information:
 - Deployment Backup OCID
 - Compartment
 - When the backup was created
 - When the backup was last updated
 - When the backup started
 - When the backup finished
 - Backup type
 - Backup size
 - Source type (manual or automatic)
 - Backup source
 - Backup location
 - Object name in Oracle Object Store
 - GoldenGate version
- View Work Requests

Copy deployment backup

To copy a deployment backup:

1. Choose one of the following options:
 - a. On the Deployments page, select the deployment you want to backup. Under Resources, select **Deployment backups**. In the Deployment backups page, select **Copy** from the Action (three dots) menu.
 - b. On the Deployment Backups page, select **Copy** for a deployment backup in the Resources section of a deployment's details page.
 - c. On the Deployment Backups page, select a deployment backup. In the Deployment backup details page, click **Copy**.
2. On the Copy backup panel, select a compartment from the dropdown and an Object storage bucket from the dropdown.
3. Click **Copy**.

Your copy appears in the list of deployment backups.

**Note:**

If you want a local copy of your deployment backup, you can download it from your Object Storage bucket.

Restoring a deployment from a backup

You can use a backup to restore a deployment to the state when the backup was created.

To restore a deployment backup:

1. On the Deployment Details page, under **Resources**, click **Deployment Backups**.
2. From the list of Deployment Backups, select a backup to restore.
3. On the Deployment Backup Details page, click **Restore**.
4. In the Restore dialog, click **Restore**.

Creating a Deployment Clone

You can use deployment backups to create a clone of a deployment.

You can create a deployment clone from the Deployment Backup Details page. Click **Create Clone**, and then complete the Create Deployment Clone wizard, which is identical to the Create Deployment wizard. For more information, see [Create deployments](#).

Monitor performance

Learn about different ways to monitor OCI GoldenGate health and performance in both the Oracle Cloud console and OCI GoldenGate deployment console.

Articles in this section:

- [Monitor performance using the Oracle Cloud console](#)
- [Monitor performance using the OCI GoldenGate deployment console](#)

Monitor performance in the Oracle Cloud console

Observe and maintain the health of your OCI GoldenGate resources by regularly monitoring metrics, creating alarms, and subscribing to events to keep informed of any abnormal activity among your resources.

Metrics

In the Oracle Cloud console, you can view metrics on the deployment details page and the metrics explorer for any Extracts and Replicats marked as critical. You can also create alarms to stay informed of certain events and take action when needed.

 **Note:**

Ensure that you upgrade your deployment to the latest version to leverage all available metrics. You must also select **Critical to deployment health** in the OCI GoldenGate deployment console for Extracts and Replicats for which to view metrics.

You can view the following OCI GoldenGate metrics in the Oracle Cloud console on the deployment details page and the metrics explorer for Extracts and Replicats marked as critical:

- **CPU Utilization:** The aggregate of all Oracle Compute Units (OCPUs).
- **CPU Consumption:** The aggregate number of OCPUs consumed.
- **Memory Utilization:** The percentage of aggregate memory. Each OCPUs allocates 16 GB memory.
- **Deployment Overall Health:** The health score of the deployment, which is the aggregate health of the deployment's processes (Administration, Distribution, Receiver, and Performance Metric Services)
- **Deployment Inbound Lag:** Lag captured for Extracts designated as critical. This metric is the aggregate of all Extracts in the deployment.
- **Deployment Outbound Lag:** Lag captured for Replicats designated as critical. This metric is the aggregate of all Replicats in the deployment.
- **Swap Space Usage:** Amount of swap space, in gigabytes, used by the deployment
- **Temp Space Usage:** Amount of temporary space, in gigabytes, used by the deployment
- **File System Usage:** Amount of file system space, in gigabytes, used by the deployment
- **Extract Status:** Health percentage of an Extract process in the deployment.
- **Replicat Status:** Health percentage of a Replicat process in the deployment.
- **DistributionPathStatus:** Health percentage of a Distribution Path process in the deployment.
- **ReceiverPathStatus:** Health percentage of a Receiver Path process in the deployment.
- **ExtractLag:** Average lag, in seconds, of a Extract process in the deployment.
- **ReplicatLag:** Average lag, in seconds, of a Replicat process in the deployment.
- **DistributionPathLag:** Average lag, in seconds, of a Distribution Path process in the deployment.
- **ReceiverPathLag:** Average lag, in seconds, of a Receiver Path process in the deployment.

 **Note:**

You'll only see metrics for the processes used in your deployment. For example, if you only have an Extract and a Replicat, then you won't see Distribution or Receiver Path metrics.

To learn more about about these metrics, see the OCI GoldenGate Metrics reference and Troubleshoot using the Oracle Cloud console.

Learn more about the [Oracle Cloud Monitoring](#).

Create Alarms

For each metric on the Deployment Details page, you can create an alert to inform you when a condition is met. For example, you can create an alarm to notify you when OCPU consumption is less than 50%.

To create an alarm:

1. From the **Options** dropdown of a metric chart, select **Create an Alarm on this Query**.
2. On the **Create Alarm** page, under **Define Alarm**, add the trigger.
3. For **Alarm Settings**, complete the following fields as needed:
 - **Alarm Name:** Enter the name that serves as the title for notifications related to this alarm. Avoid entering confidential information.
 - **Alarm Severity:** Select the perceived type of response required when the alarm is in the firing state.
 - **Alarm Body:** Enter the content of the notification to deliver.
 - **Tags (optional):** Select or enter free-form tags to apply to this resource.
 - **Metric description:** The metric to evaluate for the alarm condition.
 - **Compartment:** Select the compartment that contains the resources that emit the metrics evaluated by the alarm. The selected compartment is also where the alarm is stored.
 - **Metric Namespace:** Enter the service or application emitting metrics for the resources that you want to monitor.
 - **Resource Group (optional):** Select the group that the metric belongs to.
 - **Metric Name:** Enter the name of the metric. Only one metric can be specified.
 - **Interval:** Select the aggregation window, or the frequency at which data points are aggregated.
 - **Statistic:** Select the aggregate function.
4. Confirm the values for **Metric dimensions**. Optionally, click **+ Additional dimension** to add another dimension to the alarm.
5. For **Trigger rule**, complete the **Operator**, **Value**, and **Trigger delay minutes** fields. The graph displays the boundaries for which the alarm triggers a notification.
6. For **Notifications**, complete the fields as needed:
 - For **Destination service**, select **Notifications Service**.
 - For **Compartment**, select the compartment to store the topic used for this notification.
 - For **Topic**, click **Create topic** to set up a topic and subscription protocol in the designated compartment using the designated Destination service.

- (Optional) Click **+ Additional destination service** to add another destination service.
- (Optional) Enable **Repeat Notification** and select **Notification Interval** if you want the alarm to resend notifications at the specified intervals when the alarm is in the firing state.
- (Optional) Enable **Suppress Notifications** to specify a window of time to suspend evaluations and notifications. This is useful for maintenance periods.

7. Click **Save alarm**.

For more information, see [Viewing Default Metric Charts](#).

Subscribe to Events

Events are structured messages that indicate changes in resources. Subscribing to OCI GoldenGate events enable you to keep informed of abnormal activity among your OCI GoldenGate resources, as well as when your deployment reaches its storage limit and when an upgrade is available.

Create *rules* to subscribe to these events and trigger *actions*. For example, a rule might specify `goldengate.stateneedsattention` triggers the Notifications service to send an email to your systems administrator.

For more information, see:

- GoldenGate Events, for the full list of event types that you can create rules for.
- [Events Overview](#), to learn more about events, and how to create rules, and actions.

Monitor performance using the OCI GoldenGate deployment console

Learn to use the reports and Performance Metrics service in the OCI GoldenGate deployment console to monitor the deployment's health and performance.



Note:

This information applies only to Data replication deployments.

Monitoring Oracle GoldenGate Service Performance

You can monitor the performance of GoldenGate processes within the OCI GoldenGate deployment console using the Performance Metrics Service. The metrics service collects and stores instance deployment performance results.

After you log in to the OCI GoldenGate Deployment Console, click **Performance Metrics Service**. All the OCI GoldenGate processes are shown in their current state. Select a process to view its performance metrics.

Each service provides an elaborate view of the processes, threads, trail files, database configuration, and so on, depending on the service that you are viewing. 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 services, see the following table:

Metrics Report Tab	Available with Service
--------------------	------------------------

Process Performance	<ul style="list-style-type: none"> • Administration Service • Distribution Service • Performance Metrics Service • Receiver Service • Extracts • Replicats
Thread Performance	<ul style="list-style-type: none"> • Administration Service • Distribution Service • Performance Metrics Service • Receiver Service • Extracts • Replicats
Status and Configuration	<ul style="list-style-type: none"> • Administration Service • Distribution Service • Performance Metrics Service • Receiver Service • Extracts • Replicats
Server Statistics	<ul style="list-style-type: none"> • Distribution Service • Performance Metrics Service
Trail Files	<ul style="list-style-type: none"> • Extracts • Replicats
Database Statistics	<ul style="list-style-type: none"> • Extracts • Replicats
Procedure Statistics	<ul style="list-style-type: none"> • Extracts • Replicats
Cache Statistics	Extracts
Queue Statistics	Extracts

Reviewing Messages

Messages from the services are displayed in Performance Metrics Service home page.

To review the messages sent or received, do the following:

1. In the OCI GoldenGate Deployment Console, 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.

Review Status Changes

Real-time status changes to services 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 OCI GoldenGate Deployment Console 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.

9

Upgrade

Learn how to maintain your OCI GoldenGate deployments.

Articles in this section:

- [Maintain your OCI GoldenGate deployments](#)

Maintain OCI GoldenGate deployments

Learn about GoldenGate versions, how and when to upgrade, receive or snooze notifications, and how to rollback upgrades.

About GoldenGate versions

OCI GoldenGate supports multiple concurrent versions, for example, Oracle GoldenGate, Oracle GoldenGate for Big Data, and Oracle GoldenGate for MySQL, to name a few. Refer to the versions reference for details on the appropriate version for you.

Every release has a build number. There are three types of releases:

- **Major:** The first number in build number indicates a major release. For example, **21**.
- **Bundle:** The second number indicates the bundle release. A bundle release is a set of bug fixes for a major release. For example, **21.5**
- **Minor:** The trailing numbers indicate a minor release. A minor release consists of one or more bug fixes on a bundle release. For example, **21.5.0.0.0**.

Note:

All release types can contain security fixes.

Deployments must be upgraded when a newer version is available. Depending on the type of release and whether or not it includes a security fix, you have a specific amount of time to upgrade:

Release type	Major	Bundle	Minor
Non-security fix	365 days	180 days	Not applicable
Security fix	Not applicable	14 days	14 days

If you don't upgrade manually within the given timeframe, then your deployment automatically upgrades to the latest version at the end of this timeframe.

**Note:**

You can view the date a version is supported until in the Upgrade deployment screen and Upgrades list.

Deprecation of versions

A GoldenGate version is available from the date it was released to the date it's deprecated. When a version is deprecated, you can no longer select it for deployment creation or upgrade.

Deprecation periods use the same time frames as version upgrades. For example:

- 365 days after a Major version release, the previous Major version is deprecated.
- 180 days after a Bundle version release, previous Bundle versions of the same Major release are deprecated.
- 14 days after a Security fix release (for Bundle or Minor releases), all previous versions of the same Major release are deprecated.

**Note:**

The release date may differ between regions.

Schedule upgrades

You can customize maintenance windows that define the start of the time period during which to upgrade your deployment when a new GoldenGate version is available. If you don't define a maintenance window, then OCI GoldenGate calculates the best time to upgrade the deployment based on the latest version's release date.

**Note:**

When a maintenance window is not defined, OCI GoldenGate schedules upgrades on the weekend closest to the calculated end of the auto upgrade period.

You can customize the OCI GoldenGate maintenance window when you create the deployment, or later from the deployment details page. When a new GoldenGate version is available, the deployment automatically upgrades on the exact day and time chosen. However, deployments can automatically upgrade before your chosen day and time when the new version's timeframe for manual upgrade falls outside your customized maintenance window.

For example, let's say on January 1st, 2023, you have 30 days left to upgrade your deployment, then you have until Tuesday, January 31st to manually upgrade. However, your customized maintenance window begins on Sundays at 10PM. With this maintenance window, OCI GoldenGate will upgrade your deployment on Sunday,

February 4th at 10PM, because OCI GoldenGate will always prioritize your customized window.

To edit the customized maintenance window:

1. On the deployment details page, under GoldenGate in the Deployment information area, for Maintenance, click **Edit**.
2. In the Edit maintenance parameters dialog:
 - a. (Optional) Select **Customize maintenance window** to change the following fields from their default values.
 - b. For **Major release auto-upgrade period in days**, enter a value between 0 and 365.
 - c. For **Bundle release auto-upgrade period in days**, enter a value between 0 and 180.
 - d. For **Security path auto-upgrade period in days**, enter a value between 0 and 14.
 - e. (Optional) Select **Enable interim release auto-upgrade**, and then enter a value for **Interim release auto-upgrade period in days**.
3. Click **Save changes**.

Before you upgrade

Before upgrading a deployment, ensure that you complete the following tasks and check for long running transactions.

Ensure Archive Logs are available for recovery

Use the following command in Admin Client to determine the oldest archive log that you might need to restore when Extract starts. The `Recovery Checkpoint` field shows the oldest log needed for recovery.

```
Admin Client > INFO EXTRACT group_name, SHOWCH
```

It's best to perform upgrade activities outside of the peak hours. If there are large and long running transactions, you may consider that on the source system, the new Extract might need to start processing from the normal recovery checkpoint, rather than the bounded recovery checkpoint, if the first record of the oldest open transaction at the time that you stop Extract is in a log that is not on the system.

Clear long running transactions

To clear long running transactions, you have two options:

- You can restore the archives back to, and including, the one shown in the recovery checkpoint shown with:

```
Admin Client > INFO EXTRACT
```

(If the source database supports this.)

- You can clear out the long-running transactions that apply to the Extract that you are upgrading. This can be done by skipping the transactions or by forcing them to the trail as committed transactions. Skipping a transaction may cause data loss, and forcing a

transaction to the trail may add unwanted data to the trail if the transaction is rolled back. To skip or force a transaction:

1. View open transactions:

```
Admin Client > SEND EXTRACT group_name, SHOWTRANS
```

2. Record the transaction ID of any transaction you want to clean up.
3. Clean up old transactions using `SEND EXTRACT` and either the `SKIPTRANS` option to skip a transaction, or `FORCETRANS` to force a transaction in its current state to the Trail as a committed transaction. For example:

```
Admin Client > SEND EXTRACT group_name, {SKIPTRANS | FORCETRANS  
transaction_ID [THREAD n] [FORCE]}
```

4. After you finish cleaning up long running transactions, force a Bounded Recovery checkpoint.

```
Admin Client> SEND EXTRACT group_name, BR BRCHECKPOINT IMMEDIATE
```

Synchronize Replicats

For deployments with Parallel Replicats, ensure that you synchronize Replicats before you upgrade:

```
Admin Client > SYNCHRONIZE REPLICAT group_name
```

For more information, see `SYNCHRONIZE REPLICAT` command line interface reference.

Upgrade a deployment

Ensure that you complete the Before you upgrade steps first.

When you upgrade a deployment,

- The deployment is stopped and then restarted after the upgrade completes.
- All Oracle GoldenGate processes are stopped and then restarted if they're configured to auto-start.

To upgrade a deployment:

1. On the Deployments page, select the deployment you want to upgrade.
2. On the deployment details page, from the **More actions** dropdown, select **Upgrade**.
3. In the **Upgrade deployment** panel, select an available GoldenGate version to upgrade to, and then click **Upgrade**.

Note:

Learn more about versions.

The upgrade takes a few minutes to complete. Click **Upgrades** in the **Resources** menu on the deployment details page after the upgrade completes to view the upgrade history or if an issue is encountered during the upgrade process.

Upgrade Heartbeat tables

You must also upgrade Heartbeat tables if they're used in your deployment. After you complete the upgrade, you can either:

- Run the `UPGRADE HEARTBEATTABLE` command in Admin Client to add extra columns for tables and lag views. GoldenGate uses these extra columns to track the Extract restart position. [Learn more.](#)
- In the deployment console, open the navigation menu for the Administration console, and then click Configuration. Connect to the database. When Heartbeat tables appear, select Upgrade from the Action menu.

Upgrade notifications

OCI GoldenGate sends you event notifications (`com.oraclecloud.goldengate.upgradenotification`) about upcoming upgrades periodically. [Learn more about events.](#)

Time to upgrade	< 180 days	< 30 days	< 7 days
Notification period	Every 30 days	Every 7 days	Every day



Note:

OCI GoldenGate sends notification for security fixes every day.

You can also learn about upcoming upgrades from the notification banners that appear on the deployment details page.

Snooze notifications

You can snooze, or delay, event and banner notifications. The snooze period changes as you get closer to the end of the timeframe for manual upgrade.

Time to upgrade	< 180 days	< 30 days	< 7 days
Snooze period	30 days	7 days	Not allowed



Note:

Snooze for security fixes of Minor or Bundle upgrades is not allowed.

You can cancel snooze at any time.

Rollback upgrades

You can rollback a deployment version to the previous one. Rollback is only allowed for your latest successful upgrade to its previous version.



Note:

During rollback, the deployment file system is restored to its last state before upgrading.

Rollback is not allowed when the previous version is deprecated.

To rollback an upgrade:

1. From the Deployments page, select the deployment you want to rollback.
2. On the deployment details page, under **Resources**, click **Upgrades**.
3. In the Upgrades list, locate the most recent previous version, and then from its Action menu, select **Rollback**.
4. Select the version to rollback to, and then click **Rollback**.



Note:

Refer to the versions reference for a list of available versions.

Your deployment's status changes to updating and takes a few minutes to complete.

Reschedule upgrades

You can reschedule your auto scheduled upgrades. You can reschedule the upgrade to a closer date or postpone it to a later date.



Note:

Rescheduling doesn't take maintenance windows into account. Your newly scheduled date can't be later than OCI GoldenGate's defined value. See [About GoldenGate versions](#) for more information.

To reschedule an upgrade:

1. From the Deployments page, select the deployment whose upgrade you want to reschedule.
2. On the deployment details page, under **Resources**, click **Upgrades**.
3. In the Upgrades list, locate the most recent previous version, and then from its Action menu, select **Reschedule**.
4. In the Reschedule deployment upgrade window, select the date and time to perform the upgrade.

5. Click **Save changes**.

In the Upgrades list, the Scheduled date and time reflects the changes.

Cancel upgrades

You can only cancel an auto scheduled upgrade if the target version is an interim release.



Note:

Canceling an upgrade to an interim release won't disable the auto upgrades for interim releases implicitly and must be disabled manually. You can disable interim release auto upgrade from the deployment details page. See Schedule upgrades for more information.

To cancel an upgrade:

1. From the Deployments page, select the deployment whose upgrade you want to cancel.
2. On the deployment details page, under **Resources**, click **Upgrades**.
3. In the Upgrades list, locate the most recent previous version, and then from its Action menu, select **Cancel**.
4. In the Cancel deployment upgrade window, click **Cancel deployment upgrade**.

The upgrade's status is updated to Cancelled.

10

Secure

Learn the key areas to routinely audit to ensure your data remains secure.

Articles in this section:

- [Securing OCI GoldenGate](#)

Securing OCI GoldenGate

Oracle Cloud Infrastructure GoldenGate provides a secure and easy to use data replication solution in accordance with industry-leading security best practices.

Responsibilities

To use OCI GoldenGate securely, learn about your security and compliance responsibilities.

In general, Oracle provides security of cloud infrastructure and operations, such as cloud operator access controls and infrastructure security patching. You are responsible for securely configuring your cloud resources. Security in the cloud is a shared responsibility between you and Oracle.

Oracle is responsible for the following security requirements:

- **Physical security:** Oracle is responsible for protecting the global infrastructure that runs all services offered in Oracle Cloud Infrastructure. This infrastructure consists of the hardware, software, networking, and facilities that run Oracle Cloud Infrastructure services.
- **Encryption and confidentiality:** Encryption keys and secrets are stored in wallets and vaults to protect your data and connect to secured resources.
- **Network traffic:** Encrypted access to the OCI GoldenGate deployment console is enabled over SSL on port 443 only. By default, only access to the OCI GoldenGate deployment console is only available from an OCI private endpoint from the customer's private network. Public endpoints can be configured allowing encrypted public access to the GoldenGate Deployment Console over SSL on port 443.

Your security responsibilities include the following:

- **Access control:** Limit privileges as much as possible. Users should be given only the access necessary to perform their work.
- **OCI GoldenGate deployment console account management:** Access to the OCI GoldenGate deployment console is managed directly within the Oracle Cloud console. Accounts and permissions are managed directly in the OCI GoldenGate deployment console. Learn more about deployment users.
- **Network traffic:** Connections specify network connectivity to sources and targets. When you create a connection, you can configure SSL parameters to ensure the connection can be secure and encrypted. Learn more about connections.

- **Network encryption:** By default, all network connectivity to OCI GoldenGate is encrypted over SSL with Oracle provided certificates. Ensure that any certificate or encryption keys you provide are current and valid.
- **Audit of security events:** The OCI GoldenGate deployment console logs security events. You can access and review this log from the OCI GoldenGate deployment backup. Ensure that you monitor this log regularly. Learn more about deployment backups.
- **Patching:** Ensure that OCI GoldenGate deployments are up to date. Updates are released monthly, and you must upgrade to the latest deployment patch level as soon as possible to prevent vulnerabilities. Learn more about patching deployments.
- **Audit of remote access over Load Balancer or Bastion:** Ensure auditing of any remote access that is not directly to OCI GoldenGate is enabled and configured appropriately. [Learn more](#).

Recommendations

- Create additional OCI GoldenGate deployment console users with roles other than Security.
- Assign the minimum necessary privilege access for IAM users and groups to resource types in `goldengate-family`.
- To minimize loss of data from inadvertent deletes by an unauthorized user or malicious deletes, Oracle recommends giving the `GOLDENGATE_DEPLOYMENT_DELETE` and `GOLDENGATE_CONNECTION_DELETE` permissions to the minimum possible set of IAM users and groups. Give these permissions only to tenancy and compartment administrators.
- OCI GoldenGate only needs `USE` level access to capture data from connections.

Examples

Prevent the deletion of deployments

Create this policy to allow the group `ggs-users` to perform all actions on deployments, except deleting them:

```
Allow group ggs-users to manage goldengate-family in tenancy where  
request.permission!='GOLDENGATE_DEPLOYMENT_DELETE'
```

See Oracle Cloud Infrastructure GoldenGate Policies for more information about creating policies.

11

Troubleshoot

Should you encounter issues, learn how to troubleshoot using the resources available to you.

Articles in this section:

- [Troubleshoot using the Oracle Cloud console](#)
- [Troubleshoot using logs](#)
- [Troubleshoot using the OCI GoldenGate deployment console](#)
- [Troubleshoot memory consumption issues](#)
- [Troubleshoot connectivity issues](#)
- [Get help](#)

Troubleshoot using the Oracle Cloud console

Learn to troubleshoot OCI GoldenGate using metrics found in the Oracle Cloud console.

Deployment Information

You can use the following information in the Deployment Information tab to help you troubleshoot:

- **OCPU Count:** The base number of Oracle Compute Units (OCPUs) the OCI GoldenGate deployment has available to consume, without auto scaling. This is also the minimum meter for OCI GoldenGate.
- **Auto Scaling:** When enabled, the OCI GoldenGate deployment can scale up to three times the OCPU Count value.
- **Public IP:** If public endpoint was enabled when the OCI GoldenGate deployment was created, then the public IP is shown.
- **Private IP:** The private IP that can be accessed from your (the customer's) subnet.
- **Console URL:** The FQDN that can be used to access the OCI GoldenGate Deployment Console, over a public or private network. If private, then the console URL must be accessed from the private network.
- **OCID:** The deployment's Oracle Cloud Identifier (OCID) that is required for opening a service request (SR) with Oracle Support.

Metrics

Metrics are collected every five minutes for each deployment. The data produced can help you troubleshoot issues that you may encounter.

- **CPU Utilization:** The aggregate of all OCPUs. For example, if you specify 3 as the OCPU Count and enable Auto Scaling when you create the deployment, then the total

OCPUs that can be used is 9. When the utilization is above 33.333%, it means 33.333% of 9 OCPUs.

- **CPU Consumption:** The aggregate number of OCPUs consumed. For example, when OCPU Utilization is greater than 33.333% of 9 OCPUs, you are billed for the integer value over 33.333%, which is 4 OCPUs. When Auto Scaling is not enabled, you're billed for the base number of OCPUs.
- **Memory Utilization:** The percentage of aggregated memory. Each OCPU allocates 16 GB memory.
- **Deployment Overall Health:** Each deployment has a health score, which is the aggregate health of the underlying OCI GoldenGate deployment processes: Administration Service, Distribution Service, Receiver Service, and Performance Metrics Service.
 - Healthy = 1
 - Unhealthy = 2For example, if two of the four processes are healthy, then the health score is 50%.

 **Note:**

When you add a subprocess, such as an Extract or Distribution Path, you can designate it as **Critical to Deployment Health**. If the subprocess is stopped, then the Administration Service is deemed unhealthy.

- **Deployment Inbound Lag:** Lag is captured for Extracts that are designated as critical. This metric is aggregated across all critical Extracts.
- **Deployment Outbound Lag:** Lab is captures for Replicats that are designated as critical. This metric is aggregated across all critical Replicats.

For more information, see Metrics.

Example: Troubleshooting deployment health

This example shows you how to troubleshoot when deployment health is not 100%.

To troubleshoot deployment health in the OCI GoldenGate Deployment Console:

1. Create alarms to evaluate Deployment Health.
You'll receive notifications when Deployment Health is less than 100%.
2. Launch the OCI GoldenGate Deployment Console from the Deployment Details page and sign in.
3. In the OCI GoldenGate Deployment Console, click **Performance Metrics Service** and review the status of each process.



If a subprocess like Extract or Replicat is stopped, it directly affects the Administration Service health, giving a health score of 0 (unhealthy). Therefore the overall deployment health is 75% because only three of the four processes are healthy.

Log files are also available for each process. For more information about how to troubleshoot using the OCI GoldenGate Deployment Console log files, see [Troubleshoot using the deployment console](#).

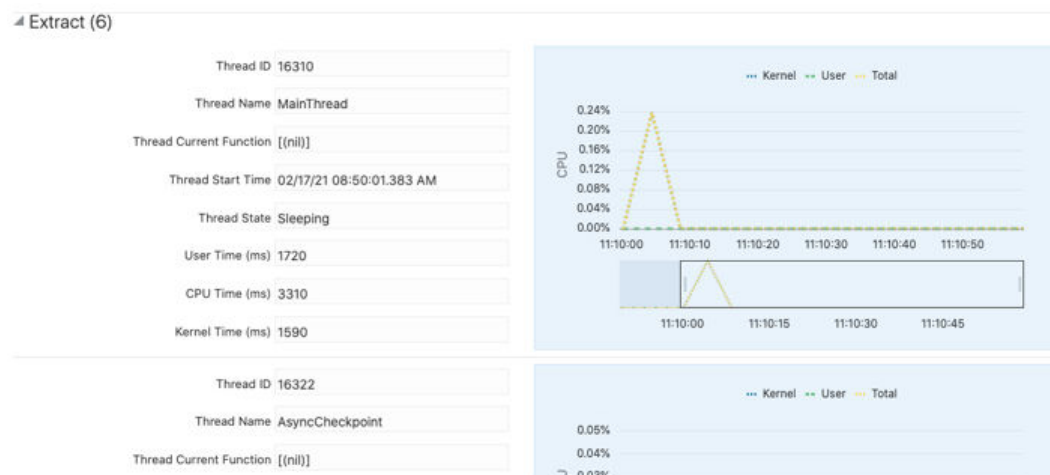
Example: Troubleshooting OCPU Utilization

This examples shows you how to troubleshoot when OCPU Utilization is greater than 90%.

Extracts and Replicats consume OCPU cycles as they replicate data. Parallel Replicats create many applicer processes for each Replicat process. After reviewing the performance metrics in the OCI GoldenGate Deployment Console, additional OCPUs may need to be added to the OCI GoldenGate deployment, or enable Auto Scaling if it's not enabled.

To troubleshoot OCPU Utilization:

1. Launch the OCI GoldenGate Deployment Console and sign in.
2. Click **Performance Metrics Service**.
3. Click each process to review its details, and then click **Thread Performance** to see the status of each thread in that process.



This information can be used to troubleshoot each process, including CPU consumption of each thread.

Troubleshoot using logs

Learn to troubleshoot common issues using OCI logging and logs found in your deployment backups, and collect diagnostics from your OCI GoldenGate deployment.

Process and error logs

You can integrate OCI Logging with OCI GoldenGate to enable, manage, and search GoldenGate process and error logs. Integration with the Logging service is optional.

- **Error logs:** Contain content from the `ggserr.log` file that logs processing events, messages, errors, and warnings generated by GoldenGate.
- **Process logs:** Contain content from several log files for Administration Service, Distribution Service, Performance Metrics Service, Receiver Service, and Extract/Replicat Events.

Before you start using OCI Logging:

- Get familiar with basic concepts and terminology used in the OCI Logging service. See [Logging Overview](#) in the OCI Logging documentation.
- Create a group to manage access to log groups and log content. See [Working with Groups](#) in the OCI IAM documentation.
- Add policies to enable public logging for OCI GoldenGate.

Note:

You manage the lifecycle of the logs including logs that are automatically created for you by OCI GoldenGate. Logs aren't deleted when the job and job runs are deleted. Learn more about [Observability and management pricing](#).

Add policies to use OCI Logging with OCI GoldenGate

To enable service logs, you must grant your user manage access on the log group, and access to the resource. Logs and log groups use the `log-group` resource-type, but to search the contents of logs, you must use the `log-content` resource-type. Add the following policies:

```
allow group <group-name> to manage log-groups in compartment
<compartment-name>
allow group <group-name> to manage log-content in compartment
<compartment-name>
```

Learn more about policies.

Enable Logging using the Oracle Cloud console

1. On the **Deployments** page, select the deployment you want to enable logging for.
2. On the **Deployment details** page, under resources select **Logs**.
3. In the **Logs** table, under the Enable log column, select the **Not enabled** toggle to enable either Process logs or Error logs.
4. In the **Enable log** panel, from the Compartment dropdown, select a compartment.

5. For Log group, you can:
 - Select a group from the dropdown
 - Create a new group
 - Leave it blank, and a default group is automatically assigned
6. For Log name, enter a name.
7. For Log retention, select the number of months from the dropdown.
8. Click **Enable log**.

Wait for the status to become Active. After it is active, the Diagnostic Logs feature is turned on for the process and error logs. The process and error logs are technically 'Service logs' as they come from an Oracle Cloud Infrastructure native service like OCI GoldenGate.

Enable OCI Logging using CLI

1. Create Log group

Request

```
oci --profile <profile_name> logging log-group create --compartment-id
<compartment_oci>
    --display-name <display_name>
```

Response

```
{ "opc-work-request-id": "<log-workrequest-oci>" }
```

2. Get Log Group ID

Request

```
oci --profile <profile_name> logging work-request get --work-request-id
<log-workrequest-oci>
```

Response

```
{
  "data": {
    "compartment-id": "<compartment_oci>",
    "id": "<log-workrequest-oci >",
    "operation-type": "CREATE_LOG_GROUP",
    "percent-complete": 100.0,
    "resources": [
      {
        "action-type": "CREATED",
        "entity-type": "loggroup",
        "entity-uri": "/logGroups/<log-group-oci>",
        "identifier": "<log-group-oci>"
      }
    ],
    "status": "SUCCEEDED",
    "time-accepted": "2023-05-09T05:57:09.641000+00:00",
    "time-finished": "2023-05-09T05:57:09.641000+00:00",
  }
}
```

```

        "time-started": "2023-05-09T05:57:09.641000+00:00"
      }
    }
  }
}

```

3. Create/Enable Log object

Request

```

oci --profile <profile-name> logging log create --display-name
<display-name> --log-group-id
    <log-group-ocid> --log-type SERVICE/CUSTOM --is-enabled true
--configuration file://path_to_json_file

```

JSON request payload

```

{
  "compartment-id": "<compartment-ocid>",
  "source":
    {
      "resource": "<GoldenGate-deployment-ocid>",
      "service": "goldengate",
      "source-type": "OCISERVICE",
      "category": "error_logs/process_logs"
    }
}

```

Response

```

Response
{
  "opc-work-request-id": "<log-workrequest-ocid>"
}

```

4. Disable OCI logging

Request

```

oci --profile <profile_name> logging log update --is-enabled false
--log-group-id <log-group-ocid> --log-id <log-ocid>

```

Response

```

{
  "opc-work-request-id": "<log-workrequest-ocid >"
}

```

5. Search OCI Logs

Request

```

oci --profile <profile_name> logging-search search-logs
    --search-query "search \"<compartment_ocid>/<log-group-ocid>/
<log-ocid>\" --time-start '2023-05-09 09:54' --time-end
'2023-05-09 09:55'"

```

Response

```

{
  "data": {
    "fields": null,
    "results": [
      {
        "data": {
          "datetime": 1683626095205,
          "logContent": {
            "data": {
              "level": "INFO",
              "message": "Executing command '{\n  \"name\":
\nreport\", \n  \"reportType\": \"lag\", \n  \"thresholds\": [\n
\n  \"type\": \"info\", \n  \"units\":
\nseconds\", \n  \"value\": 0\n  ]
\n}'",
              "processName": "adminsrvr",
              "resourceId": "<GoldenGate-deployment-
ocid>"
            },
            "id": "20230509095455.2051683626095",
            "oracle": {
              "compartmentid": "<compartment-
ocid>",
              "ingestedtime":
"2023-05-09T09:55:05.889Z",
              "loggroupid": "<log-group-ocid>",
              "logid": "<log-ocid>",
              "tenantid": "<tenant-ocid>"
            },
            "source": "<GoldenGate-deployment-ocid>",
            "specversion": "1.0",
            "time": "2023-05-09T09:54:55.205Z",
            "type":
"com.oraclecloud.goldengate.deployment.process_logs"
          }
        }
      }
    ],
    "summary": {
      "field-count": null,
      "result-count": 1
    }
  }
}

```

6. Get Log Status**Request**

```

oci --profile <profile_name> logging log get --log-group-id <log-group-
ocid> --log-id
<log-ocid>

```

Response

```
{
  "data": {
    "compartment-id": "<compartment-ocid>",
    "configuration": {
      "archiving": {
        "is-enabled": false
      },
      "compartment-id": "<compartment-ocid>",
      "source": {
        "category": "process_logs",
        "parameters": {},
        "resource": "<GoldenGate-deployment-ocid>",
        "service": "goldengate",
        "source-type": "OCISERVICE"
      }
    },
    "defined-tags": {
      "Oracle-Tags": {
        "CreatedBy": "<creator_email_id>",
        "CreatedOn": "2023-05-09T06:24:30.279Z"
      }
    },
    "display-name": "<display_name>",
    "freeform-tags": {},
    "id": "<log-ocid>",
    "is-enabled": false,
    "lifecycle-state": "INACTIVE",
    "log-group-id": "<log-group-ocid>",
    "log-type": "SERVICE",
    "retention-duration": 30,
    "tenancy-id": "<tenancy-ocid>",
    "time-created": "2023-05-09T06:24:30.452000+00:00",
    "time-last-modified": "2023-05-09T06:30:17.345000+00:00"
  },
  "etag": "cb5bb295-9954-4949-b33f-67d6da50f83f"
}
```

7. Delete Log**Request**

```
oci --profile <profile_name> logging log delete --log-group-id <log-group-ocid> --log-id <log-ocid>
```

Response

```
{
  "opc-work-request-id": "<log-workrequest-ocid >"
}
```

8. Delete Log group

Request

```
oci --profile GGSTEST logging log-group delete --log-group-id <log-group-ocid>
```

Response

```
{
  "opc-work-request-id": "<log-workrequest-ocid >"
}
```

Enable Logging for OCI GoldenGate in the OCI Logging service

Create a log group and configure a [service log](#) in the Logging service if you don't have one already:

1. Open the navigation menu and click **Observability & Management**. Under **Logging**, click **Log Groups**.
2. Choose a compartment you have permission to work in and click **Create Log Group**. The **Create Log Group** panel is displayed.
3. Complete the following:
 - **Compartment:** The compartment in which you want to create the log group. This field is pre-filled based on your compartment choice.
 - **Name:** A name for this log group. The first character of a log group name must be a letter. For more, see [Log and Log Group Names](#). Avoid entering confidential information.
 - **Description:** A friendly description.
 - (Optional) Enter tagging information.
4. Click **Create**.
The log group detail page is then displayed.
5. Click **Logs**.
6. Click **Enable service log**.
7. Under **Select Resource**:
 - a. For **Resource Compartment**, select the resource's compartment.
 - b. For **Service**, select **GoldenGate** from the dropdown.
 - c. For **Resource**, select a deployment from the dropdown.
8. Under **Configure log**:
 - a. For **Log Category**, select a value from the dropdown.
 - b. For **Log Name**, enter a name.
9. Select **Show Advanced Options**, and under **Log Location**:
 - a. For **Compartment**, select the compartment in which to save the log.
 - b. For **Log Group**, select a log group from the dropdown or create new group.
10. For **Log Retention**, select the number of months for which to retain the logs.

11. Click **Enable log**.

Troubleshoot health issues using deployment backups

You can manually backup an OCI GoldenGate deployment to Oracle Object Storage, and then download the backup locally.

To learn how to create a manual back up, see [managing deployment backups](#).

The manual deployment backup contains the full GoldenGate deployment directory structure and files, including log and trail files. The directories and files that are helpful for troubleshooting include:

- /etc: configuration
 - /etc/ogg: parameter files
- /var: log files, checkpoint, trail files, and so on
 - /var/checkpt: checkpoint
 - /var/data: trail files
 - /var/report: report files
 - /var/log: log files
 - * ER-events.log
 - * ggserr.log
 - * restapi.log
 - * adminsvr.log
 - * pmsvr.log
 - * distsvr-stdout.log
 - * recsvr-stdout.log
 - * distsvr.log
 - * recsvr.log
 - * extract.log
 - * replicat.log

Collect diagnostics

Collect diagnostics to analyze or share information about your OCI GoldenGate deployment. The information collected can be shared with My Oracle Support if you encounter any issues.



Note:

This feature applies only to Data replication deployments.

Before you collect diagnostics, ensure that you [create an OCI Object Storage bucket](#).

To collect deployment diagnostics:

1. On the Deployments page, select the deployment for which to collect diagnostics.
2. On the deployment's Details page, from the More actions menu, select **Collect diagnostics**.
3. In the Collect diagnostics panel, complete the following fields, and then click **Collect diagnostics**:
 - a. From the **Bucket** dropdown, select the bucket in which to save the diagnostics file. If you want to select a bucket in a different compartment, click **Change Compartment**.
 - b. For **Diagnostics name prefix**, enter a short name or a few characters to prefix the diagnostics file name.
 - c. (Optional) Select the Start date from which to collect system logs.
 - d. (Optional) Select the End date to which to collect system logs.

 **Note:**

Diagnostics collected for the OCI GoldenGate deployment contains GoldenGate logs for the entire lifespan of the deployment and are independent of the start and end dates selected.

After you click Collect diagnostics, a new field named **Diagnostics** displays under the GoldenGate section of the Deployment information card. It can take several minutes for the diagnostics zip file to become available to download. When it is available, a **Download** link appears.

Troubleshoot using the OCI GoldenGate deployment console

You can use log files in the OCI GoldenGate deployment console to gain more insight on issues you may encounter.

 **Note:**

This information applies only to Data replication deployments.

Depending on the type of issue, you can start in one of two places:

- **Deployment Health issues:** For health issues not caused by user action, such as stopping a critical Extract. Start troubleshooting using logs generated by the OCI GoldenGate deployment console when you create a manual backup of the deployment. You can download the backup locally, and then view logs from the backup.
- **Performance issues:** You can drill down to additional details using the Diagnosis tab, Debug log, Critical Events logs, or the Performance Metrics Server to see GoldenGate metrics in real time.

Troubleshoot using the Administration Service

In the Administration Service area, you can use the following to help you troubleshoot:

- Extract and Replicat processes:
 - Add, alter, and delete
 - Register and unregister
 - Start and stop
 - Review process information, statistics, and status, including LAG and checkpoints
 - Retrieve the report and discard files
- Log files
 - Diagnosis for Critical Events
 - Debug Log for additional debug information
- Configuration (parameter) files
- Checkpoint, trace, and heartbeat tables
- Supplemental logging for procedural replication, schema, and tables
- Tasks, both custom and standard, such as auto-restart and purge trails
- Credential stores
- Add users and assign their roles

Troubleshoot using the Distribution Service

In the Distribution Service area, you can use the following to help you troubleshoot:

- Distribution Path processes:
 - Add, alter, and delete
 - Reposition path, change filtering
 - Monitor data statistics outgoing path
- Log files
 - Diagnosis for Critical Events
 - Debug Log for additional debug information

Troubleshoot using the Receiver Service

In the Receiver Service area, you can use the following to help you troubleshoot:

- Receiver Path processes:
 - Add, alter, and delete
 - Reposition path, change filtering
 - Monitor data statistics incoming path
 - Monitor network/file IO statistics incoming path

- Log files
 - Diagnosis for Critical Events
 - Debug Log for additional debug information

Troubleshoot using the Performance Metrics Service

In the Performance Metrics Service area, you can monitor Oracle GoldenGate deployment processes to help you troubleshoot:

- Administration Service
 - Extracts
 - Replicats
- Distribution Service
- Receiver Service
- Performance Metrics Service
- Log files
 - Diagnosis for Critical Events
 - Debug log for additional debug information
- Monitor system Oracle GoldenGate deployment system wide logs
- Monitor system Oracle GoldenGate deployment system status changes

Troubleshoot memory consumption issues

If your Oracle GoldenGate Extracts or Replicats run out of memory when running long transactions, then you must configure Cache Manager properly for each process.

The default cache size setting of 64 GB is fine in most cases. An issue arises when you process extremely large transactions or your deployment shape is small relative to the transaction size you're processing.

Depending on the shape of the OCI GoldenGate deployment, 16 gigabytes (GB) of physical memory is allocated per OCPU. If autoscale was enabled, then the total physical memory can be up to 3 times more.

For example, an OCI GoldenGate deployment with 4 OCPUs (4 x 16 GB = 64 GB) and autoscale enabled (64 GB x 3) has 192 GB of physical memory available.

By default, each OCI GoldenGate deployment is allocated 256 GB of swap space. Therefore, your total virtual memory for the deployment is 192 GB + 256 GB = 448 GB. Considering that all GoldenGate processes in the deployment share this virtual memory, Oracle recommends that you set `CACHEMGR CASHESIZE` to a safe range, such as 60% of the total virtual memory, or 268 GB. Once the Extract or Replicat reaches 268 GB virtual memory usage, overflow transactions are written to disk, which defaults to `/u03/temp`. This default directory has 256 GB available. If that's still not enough space, then you can add a secondary directory. Oracle recommends `/u02/Deployment/var/lib/cachemanager`, which is unlimited.

```
CACHEMGR CASHESIZE 268G CACHEDIRECTORY /u03/temp, CACHEDIRECTORY /u02/
Deployment/var/lib/cachemanager
```

If you have multiple Extracts and Replicats that process large transactions in the same deployment, then you can choose from the following options:

- Increase the deployment's OCPUs so that you have more physical memory and results in more virtual memory for your processes to use.
- Decrease the `CACHEMGR CACHESIZE` setting so that each process uses less virtual memory and force the overflow to use the disk directories. This lets you keep the current deployment shape without the increase cost. However, your large transaction processing will be slightly slower as it uses disk instead of virtual memory to cache the pending transactions.

You can use the parameters and recommendations mentioned here to help fine tune your environment settings to your transaction size.

Troubleshoot connectivity issues with Oracle Database

Learn to troubleshoot connectivity issues in OCI GoldenGate for on-premises databases or third-party cloud databases.

Connectivity issues are among the most common errors encountered in OCI GoldenGate. If you experience an issue connecting to your on-premises database or a database running in a third-party cloud, the following tips may help you troubleshoot:

1. Ensure you review the instructions to [Create a connection to Oracle Database](#).
2. In the Network connectivity section, if you selected **Dedicated endpoint** for Traffic routing method:
 - a. If you're connecting to a Real-Application Cluster (RAC) database using [Single Client Access Name \(SCAN\)](#), then you must select **Redirect** (Session mode), and then enter the SCAN listener FQDN for **RAC node IP**.
 - b. If you're NOT connecting to a RAC database, then you can enter any FQDN (fully qualified database name) into this field. It doesn't have to be the name of a database. It does, however, have to be unique within the Compartment and in a valid format. For example, `somehost.example.com`.
 - c. You must refer to the FQDN entered in the **Database connection string** field, for example, `somehost.example.com:1521/my servicenameexample`

Internally, the FQDN value you enter gets mapped within the OCI GoldenGate service tenancy to the internal private endpoint IP address corresponding to the actual database's private IP address.

Note:

If you don't refer to the FQDN provided in the **Database connection string**, then OCI GoldenGate won't be able to connect to the database, and you could encounter the following error messages:

- Error - OGG-08110 Login failed. OCI Error ORA (status = 12170-ORA-12170: TNS:Connect timeout occurred)
- Error - OGG-08110 Login failed. OCI Error ORA (status = 12154-ORA-12154: TNS:could not resolve the connect identifier specified)

3. You can connect to databases outside of Oracle Cloud Infrastructure (OCI) using the private IP of the database.
 - a. Ensure that you select **Dedicated endpoint**, which allows connectivity from OCI GoldenGate to the database node's private IP.
 - b. Select the appropriate **Session mode**. Choose **Direct** to use the local listener running on a single database node, or **Redirect** to use the SCAN listener used in RAC deployments.
 - c. Select a subnet that has access to the **Database node IP** that you provided (optional for RAC databases). OCI GoldenGate creates a reverse connection private endpoint in that subnet to access your database.

 **Note:**

If you're trying to connect to an on-premises database or a database running in a third-party cloud, you must ensure that you configure OCI Networking to enable access from the subnet you provided into the network where the database resides. OCI offers different ways to achieve this, including [FastConnect](#) or [IPSec VPN](#).

- d. You must create security rules that allow ingress from the subnet provided for port 1521, or whatever port you're trying to access. If you don't create this security rule, then OCI blocks traffic to this port.
4. After creating the connection, review its details. Observe the **Ingress IPs** field. It may contain one or more IP addresses. The connection to the database from OCI GoldenGate originates from one of these IP addresses. You must ensure that the appropriate subnet security rules are in place to allow connectivity from these IP addresses into the database node's private IP.

Get help

You can raise a service request with **My Oracle Support** if you need help to resolve issues when working with Oracle Cloud Infrastructure GoldenGate.

Submitting a Service Request

My Oracle Support is a customer portal that offers product services through various support tools and contains a repository of useful information, where you can find solution to your issue. You can raise a service request using this application through one of the following two interfaces:

1. **My Oracle Support**
2. **Cloud Support**

You must meet the following prerequisites to create a service request:

- You must have a Support Identifier which verifies your eligibility for Support services.
 - You must have an account at **My Oracle Support**
1. Access **My Oracle Support** at <https://support.oracle.com>.

You can choose to create a service request either from **My Oracle Support** interface or from **Cloud Support** interface by using the switch toggle button on the top-right of the window.

2. Perform the following steps to create a service request from **My Oracle Support** interface:
 - a. Click **Create Technical SR** on the Service Requests tab.
 - b. Enter the **Problem Summary**.
 - c. Enter the **Problem Description**.

 **Note:**

It is important to provide your **home region**, **tenancy name** and **database name** along with your problem details. You can optionally provide your user-friendly display name.

- d. Enter the **Error Codes**.
- e. Select the **Cloud tab** under "Where is the Problem".
- f. Specify **GoldenGate Service** in the Service Type field.

 **Note:**

DO NOT select the legacy **Oracle GoldenGate Cloud Service**. Doing so will add an additional day of processing as your SR gets rerouted to the proper channels.

- g. Select a **Problem Type** and provide the **Support Identifier** details.
 - h. Click **Next** until you have provided all the mandatory information.
 - i. Click **Submit**.
Your service request is created.
3. Perform the following steps to create a service request from **Cloud Support** interface:
 - a. Click **Create Technical SR** on the Service Requests tab.
 - b. Follow through sub-steps **2.f** to **2.i** in the preceding step.
Your service request is created.

If you are asked to provide log files, follow the steps outlined in Collect Diagnostics.

12

Reference

Reference information in this section includes details about service events, metrics, and policies.

Articles in this section:

- [Frequently asked questions](#)
- [Oracle Cloud Infrastructure GoldenGate Versions](#)
- [Oracle Cloud Infrastructure GoldenGate Events](#)
- [Oracle Cloud Infrastructure GoldenGate Metrics](#)
- [Oracle Cloud Infrastructure GoldenGate Policies](#)
- [Known issues in OCI GoldenGate](#)

Frequently asked questions

Find the answer to some commonly asked questions about Oracle Cloud Infrastructure GoldenGate.

About Region Availability

In what regions is OCI GoldenGate available?

Visit [Cloud Regions](#) to see where OCI GoldenGate is available.

About Host Names and IPs

Can I use Oracle Single Client Access Name (SCAN) host names and IP addresses with OCI GoldenGate?

Yes. You can register Real Application Cluster (RAC) databases using SCAN IP addresses. See [Create a connection to Autonomous Database](#) or [Create a connection to Oracle Database](#).

Does SCAN Proxy support TLS?

TLS support for SCAN Proxy is not yet enabled. You can connect to a RAC database using the database node (SCAN) IP.

How do I connect to the OCI GoldenGate Deployment Console using a private IP?

You can use the OCI GoldenGate deployment's private URL from a Compute instance in the same VCN, or through a bastion host. See [Connect to OCI GoldenGate Using a Private IP](#).

About Ports and Protocols

What ports and protocols are available for communication to the OCI GoldenGate deployment?

Only port 443 is open to communicate to the OCI GoldenGate deployment and SSL communication is required. Therefore, access to OCI GoldenGate deployment must be over HTTPS and WSS.

About Oracle Databases

Can I connect OCI GoldenGate to an on-premises Oracle Database?

Yes, any supported technology that you can successfully connect to from your customer tenancy can be used as a source or target within OCI GoldenGate. The proper ingress and egress rules must exist to allow a successful database connection.

Can I connect to an Oracle Database that's within a third-party cloud?

Yes, any supported technology that you can successfully connect to from your customer tenancy can be used as a source or target within OCI GoldenGate. The proper ingress and egress rules must exist to allow a successful database connection.

About Oracle Autonomous Databases

Can I capture and deliver data to and from Oracle Autonomous Databases?

Yes, Oracle Autonomous Transaction Processing (ATP) and Oracle Autonomous Data Warehouse (ADW) on Exadata Infrastructure are supported for capture and delivery with OCI GoldenGate.

About APIs and SDKs

Can I execute REST calls directly to the OCI GoldenGate Deployment Console?

Yes, REST calls are fully supported. [Learn more](#).

For information about SDKs, see [Software Development Kits](#).

Oracle Cloud Infrastructure GoldenGate versions

Learn about the different OCI GoldenGate release versions and what's included in each before you upgrade.

You can subscribe to upgrade events to stay informed of the availability of new versions. Learn more about managing deployment upgrades.

About GoldenGate versions

OCI GoldenGate supports multiple concurrent versions, for example, Oracle GoldenGate, Oracle GoldenGate for Big Data, and Oracle GoldenGate for MySQL, to

name a few. Refer to the versions reference for details on the appropriate version for you.

Every release has a build number. There are three types of releases:

- **Major:** The first number in build number indicates a major release. For example, **21**.
- **Bundle:** The second number indicates the bundle release. A bundle release is a set of bug fixes for a major release. For example, **21.5**
- **Minor:** The trailing numbers indicate a minor release. A minor release consists of one or more bug fixes on a bundle release. For example, **21.5.0.0.0**.

 **Note:**

All release types can contain security fixes.

Deployments must be upgraded when a newer version is available. Depending on the type of release and whether or not it includes a security fix, you have a specific amount of time to upgrade:

Release type	Major	Bundle	Minor
Non-security fix	365 days	180 days	Not applicable
Security fix	Not applicable	14 days	14 days

If you don't upgrade manually within the given timeframe, then your deployment automatically upgrades to the latest version at the end of this timeframe.

 **Note:**

You can view the date a version is supported until in the Upgrade deployment screen and Upgrades list.

Deprecation of versions

A GoldenGate version is available from the date it was released to the date it's deprecated. When a version is deprecated, you can no longer select it for deployment creation or upgrade.

Deprecation periods use the same time frames as version upgrades. For example:

- 365 days after a Major version release, the previous Major version is deprecated.
- 180 days after a Bundle version release, previous Bundle versions of the same Major release are deprecated.
- 14 days after a Security fix release (for Bundle or Minor releases), all previous versions of the same Major release are deprecated.

 **Note:**

The release date may differ between regions.

Upgrade notifications

OCI GoldenGate sends you event notifications (`com.oraclecloud.goldengate.upgradenotification`) about upcoming upgrades periodically. [Learn more about events.](#)

Time to upgrade	< 180 days	< 30 days	< 7 days
Notification period	Every 30 days	Every 7 days	Every day



Note:

OCI GoldenGate sends notification for security fixes every day.

You can also learn about upcoming upgrades from the notification banners that appear on the deployment details page.

Oracle GoldenGate for Oracle versions

Learn about what's included in each of the latest Oracle GoldenGate versions.



Note:

- To find out when a version was released and how long you have to upgrade, refer to your deployment details page.
- To find out what new features are available, see [What's new in OCI GoldenGate](#).
- For more details on bug fixes and enhancements listed in the table below, see [Oracle GoldenGate bugs fixed and enhancements](#)

Type	Version number	Features and updates included
Minor, Security	21.12.0.0.0_231115.1435_874	Bugs fixed: <ul style="list-style-type: none"> • Bug 35145292: Oracle - INFO SCHEMATRANDATA command fails in Microservices deployment with ~81k tables • Bug 35360186: Oracle - In Oracle GoldenGate 21.9 Microservices release, only report file for replicat main process is showing in the adminsvr report file section
Minor	21.12.0.0.0_231007.0000_862	Internal updates, no customer impact.
Bundle	21.12.0.0.0_231007.0000_843	Internal updates, no customer impact.

Type	Version number	Features and updates included
Minor	21.11.0.0.0_230912.1108_82 2	Bugs fixed: <ul style="list-style-type: none"> Bug 35367875: OGG-14053 Invalid Heartbeat record encountered Bug 35661628: AdminClient Edit Utility (EDIT params {process_name}) issue Bug 35702967: PR-Non_Integrated: PR Fails with OGG-02092 Unexpected condition in function indexOutOfRangeException at line 51. Index 1078 out of range while processing very large tx Enhancements: <ul style="list-style-type: none"> Enh 35728715: Disable query to support mode on releases doesn't have query optimization
Minor	21.11.0.0.0_230714.2015_81 2	Internal updates, no updates to core GoldenGate version. No customer impact.
Minor	21.11.0.0.0_230714.2015_80 2	Internal updates, no customer impact.
Minor	21.11.0.0.0_230714.2015_78 6	Bugs fixed: <ul style="list-style-type: none"> GG-7261: Diagnostic collection fails to collect when log files are larger than 14GB
Bundle	21.11.0.0.0_230714.2015_76 7	Bugs fixed: <ul style="list-style-type: none"> Bug 35446921 - Extract Abends when binlog_row_value_options is set to partial Json Bug 35507352 - Display the LogNumber and LogPosition on the error message when the extract abends due to partial json update Bug 34262780: Oracle - Integrated Replicat fails with error "OGG-00664 OCI error ORA-01406" Bug 34277295: Oracle - Integrated Extract fails with error OGG-01112 in DDLEXT_process_extract for GRANT ddl Bug 34294074: Oracle - Parallel Replicat abends intermittently with error OGG-00418 Bug 35401796: Oracle - Integrated Replicat crashes with message "*** stack smashing detected ***: /opt/ogg/classic/19/commbank04/19c/replicat terminated" Bug 35406535: Oracle - BR files not getting cleared upon BR cancel and disabled Bug 32131835: Generic - Enhance logdump to optionally disable all column data output Bug 34732293: Generic - Mapper memory leak occurs when parallel Replicat is started with ATCSN/AFTERCSN Bug 35164853: Generic - When the GLOBALS file is not present, the AllowNullableKeys command is effectively true Bug 35225748: Generic - Extract registration information getting removed from microservices metadata Bug 35350661: Generic - Modify recvsrvr /targets and /targets/path endpoints Bug 35401164: Generic - Added a new TRANLOGOPTIONS UNSUPPORTEDDLCOPTION parameter
Minor	21.10.0.0.0_230530.1017_76 6	Internal updates, no customer impact.

Type	Version number	Features and updates included
Minor	21.10.0.0.0_230530.1017_765	Bugs fixed: <ul style="list-style-type: none"> • BUG 35256956 - Replicat process not terminating after abend • BUG 35384626 - Extract keeps loading after creating it in the WEBUI
Bundle	21.10.0.0.0_230509.1021_744	Bugs fixed: <ul style="list-style-type: none"> • Bug 34791431: Exadata on Premise - Replicat inserts garbage characters in CLOB column when character set is source database is set to CESU-8 • Bug 35097023: Exadata on Premise - Parallel Replicat abends with error 1403 and skips transactions and continues progressing after two auto restart attempts from the manager • Bug 34760678: Oracle - Parallel integrated Replicat consumes higher PGA with large # of appliers leading to PGA limit breach/4036 • Bug 34851353: Oracle - Parallel Replicat is missing transactions randomly which is causing data integrity issue • Bug 34626040: Generic - Replicat abends when trying to use <code>_LOW_WATERMARK_UPDATE_TIME</code> parameter in the parameter file • Bug 35161065: Generic - Modifying distribution path <code>tcpRcvBuf</code> or <code>tcpSndBuf</code> prevents path starting and causes <code>DistSrvr</code> to fail with core dump • Bug 35229639: Generic - ADMINSTRVR Web UI: Request ER lag by report lag command • Bug 35132740 - <code>oggscs.jar</code> and <code>reverseproxysettings.jar</code> are showing as vulnerable dependency
Minor	21.9.0.0.0_230218.0903_722	Internal updates, no customer impact.
Minor	21.9.0.0.0_230218.0903_716	Internal updates, no customer impact.
Bundle	21.9.0.0.0_230218.0903_703	Bugs fixed: <ul style="list-style-type: none"> • Bug 33701569 - MA: Generic - START REPLICAT with ATCSN/AFTERCSN does not work for cross-database positioning • Bug 34406231 - Secure Receiver Server exhibits unbounded memory consumption • Bug 35055861 - DROP PARTITION skipped by OCI GGS Extract at Intermediate Database • Bug 34851353 - Parallel Replicat is missing transactions randomly, causing data integrity issue • Bug 3498303 - User with low privileges is able to access DEBUG logs
Minor	21.8.0.0.0_230102.2154_687	Bugs fixed: <ul style="list-style-type: none"> • GGS-7217: Fix Oracle Exadata connection creation, when <code>dbsystem</code> is configured in <code>vmcluster</code>

Type	Version number	Features and updates included
Minor	21.8.0.0.0_221119.1258_663	Bugs fixed: <ul style="list-style-type: none"> Bug 34738804 - Administration Service UI/API hangs intermittently. Bug 34784214 - Alias not showing under "Credential Alias" when adding a new Extract when username has '@' Bug 34764925 - Web-UI: Drop-down menu does not expand or show all deployments
Bundle, Security fix	21.8.0.0.0_221014.1235_654	Internal updates, no customer impact.
Bundle	21.8.0.0.0_221014.1235_616	Enhancements: <ul style="list-style-type: none"> Enhancement 34466644 - Support Environment Variables in Activity Logging Patterns Bugs fixed: <ul style="list-style-type: none"> Bug 34275661 - Users listed in the endpoint /services/v2/authorizations/All are not found when IDCS is configured in deployment Bug 33701569 - MA: start replicat when atscn and aftercsn doesn't work from webui/adminclient for all DBs (PostgreSQL) Bug 34061707 - Distrsrvr looping on reconnect/send logic Bug 34530777 - Replicat hanging in OCI GoldenGate environment and reporting high CPU Bug 34424244 - Critical Extract in Abended state does not change deployment health to unhealthy Bug 34350756 - WEBUI: Oracle GoldenGate webui should by default use userdalias only and not the source DB for all DBs

Oracle GoldenGate for Big Data versions

Learn about what's included in each of the latest Oracle GoldenGate for Big Data versions.

- To find out when a version was released and how long you have to upgrade, refer to your deployment details page.
- To find out what new features are available, see [What's new in OCI GoldenGate](#).
- For more details on bug fixes and enhancements, see [Oracle GoldenGate for Big Data Release Notes](#)

Type	Version number	Features and updates included
Minor	21.12.0.0.0_231018.0822_874	Internal updates, no customer impact
Bundle	21.12.0.0.0_231018.0822_862	See Oracle GoldenGate for Big Data Release Notes for the full list of bug fixes and enhancements.
Minor	21.11.0.0.0_230920.0759_841	Internal updates, no customer impact
Minor	21.11.0.0.0_230829.0229_822	Internal updates, no customer impact
Bundle	21.11.0.0.0_230829.0229_812	Internal updates, no customer impact.

Type	Version number	Features and updates included
Minor	21.10.0.0.0_230713.1632_80 2	Internal updates, no customer impact.
Minor	21.10.0.0.0_230713.1632_78 6	Bugs fixed: <ul style="list-style-type: none"> • GGS-7261: Diagnostic collection fails to collect when log files are larger than 14GB • Bug 34948936: Generic - Memory leak in OGGBD replicat processes
Minor	21.9.0.0.0_230412.2026_766	Internal updates, no customer impact.
Minor	21.9.0.0.0_230412.2026_765	Internal updates, no customer impact.
Minor	21.9.0.0.0_230412.2026_744	Internal updates, no customer impact.
Minor	21.9.0.0.0_230412.2026_723	Bugs fixed: <ul style="list-style-type: none"> • Bug 35209917 - Snowflake key pair replicat log file shows password on java exception-negative test • Bug 35187763 - Snowflake key pair replicat log file displays password
Bundle	21.9.0.0.0_230324.1300_716	Internal updates, no customer impact.
Minor	21.7.0.0.0_230111.1953_703	Internal updates, no customer impact.
Minor	21.7.0.0.0_230111.1953_687	Bugs fixed: <ul style="list-style-type: none"> • Bug 34488932 - PMSRVR data store getting corrupted
Minor	21.7.0.0.0_221110.2013_663	Internal updates, no customer impact.
Bundle	21.7.0.0.0_221110.2013_654	Enhancements: <ul style="list-style-type: none"> • Enh 34516628 - OGGBD GGS Kafka extract Parameter to read schema registry connection id • Enh 34318121 - List connections in OCIGG Big Data UI Connections now appear in deployment console Configuration screen. • Enh 34362397 - OCI GoldenGate Big Data ADW Stage & Merge In ADW Stage & Merge Replicat properties, OCI Object Storage Connection OCID was not populated/missing.
Minor	21.6.0.0.2_220714.0903_616	Enhancements: <ul style="list-style-type: none"> • Enh 34114672 - OCIGG Big Data Kafka Extract C code changes to accept new configuration connectionId

Oracle GoldenGate for Microsoft SQL Server versions

Learn about what's included in each of the latest Oracle GoldenGate for Microsoft SQL Server versions.

 **Note:**

- To find out when a version was released and how long you have to upgrade, refer to your deployment details page.
- To find out what new features are available, see [What's new in OCI GoldenGate](#).
- For more details on bug fixes and enhancements listed in the table below, see [Oracle GoldenGate bugs fixed and enhancements](#)

Type	Version number	Features and updates included
Bundle, Security	21.12.0.0.0_231014.0710_87 4	Bugs fixed: <ul style="list-style-type: none"> • Bug 34922587: Extract is deleted when logged in to the wrong database • Bug 35651029: Modify the default "keep rule for purge change data task" option to 60 minutes
Minor	21.11.0.0.0_230722.0716_86 2	Internal updates, no customer impact.
Minor	21.11.0.0.0_230722.0716_84 1	Internal updates, no customer impact.
Minor	21.11.0.0.0_230722.0716_82 2	Bugs fixed: <ul style="list-style-type: none"> • Bug 34812191: Microsoft SQL Server - Support capture and delivery for SQL Server 2022 • Bug 35309968: Microsoft SQL Server - Credentials from Admin Service page are not editable
Minor	21.11.0.0.0_230714.2015_81 2	Internal updates, no customer impact.
Bundle	21.11.0.0.0_230714.2015_80 2	Bugs fixed: <ul style="list-style-type: none"> • GGS-7261: Diagnostic collection fails to collect when log files are larger than 14GB • Bug 35353317: MySQL - Replicat using incorrect metadata while processing DDL record when source database is case sensitive and target is case insensitive • Bug 32131835: Generic - Enhance logdump to optionally disable all column data output • Bug 34732293: Generic - Mapper memory leak occurs when parallel Replicat is started with ATCSN/AFTERCSN • Bug 35164853: Generic - When the GLOBALS file is not present, the AllowNullableKeys command is effectively true • Bug 35225748: Generic - Extract registration information getting removed from microservices metadata • Bug 35350661: Generic - Modify recvsrvr /targets and /targets/path endpoints • Bug 35401164: Generic - Added a new TRANLOGOPTIONS UNSUPPORTEDLDCROPTION parameter
Minor	21.10.0.0.0_230413.1303_78 6	Internal updates, no customer impact.
Minor	21.10.0.0.0_230413.1303_76 6	Internal updates, no customer impact.

Type	Version number	Features and updates included
Minor	21.10.0.0.0_230413.1303_765	Bugs fixed: <ul style="list-style-type: none"> • BUG 35256956 - Replicat process not terminating after abend • BUG 35384626 - Extract keeps loading after creating it in the WEBUI
Bundle	21.10.0.0.0_230413.1303_744	Bugs fixed: <ul style="list-style-type: none"> • Bug 31607593: Microsoft SQL Server - Automatically remove OracleGGExtractCheckpoint record for an Extract when that Extract is deleted • Bug 35018016: Microsoft SQL Server - Installing OGG 21.4 Microservice for SQLServer report missing vcredist program • Bug 34626040: Generic - Replicat abends when trying to use _LOW_WATERMARK_UPDATE_TIME parameter in the parameter file • Bug 35161065: Generic - Modifying distribution path tcpRcvBuf or tcpSndBuf prevents path starting and causes Distsrvr to fail with core dump • Bug 35229639: Generic - ADMIN\$RVR Web UI: Request ER lag by report lag command
Minor	21.9.0.0.0_230120.0600_722	Internal updates, no customer impact.
Minor	21.9.0.0.0_230120.0600_716	Initial release version.

Oracle GoldenGate for MySQL versions

Learn about what's included in each of the latest Oracle GoldenGate for MySQL versions.

Note:

- To find out when a version was released and how long you have to upgrade, refer to your deployment details page.
- To find out what new features are available, see [What's new in OCI GoldenGate](#).
- For more details on bug fixes and enhancements, see [Oracle GoldenGate bugs fixed and enhancements](#)

Type	Version number	Features and updates included
Bundle, Security	21.12.0.0.0_231025.1441_874	Bugs fixed: <ul style="list-style-type: none"> • Bug 33652805: Support for position by GTID set • Bug 35838196: Oracle GoldenGate incorrectly writes binary value which causes unreadable target data value
Minor	21.11.0.0.0_230714.2015_862	Internal updates, no customer impact.
Minor	21.11.0.0.0_230714.2015_841	Internal updates, no customer impact.

Type	Version number	Features and updates included
Minor	21.11.0.0.0_230714.2015_82 2	Internal updates, no customer impact.
Minor	21.11.0.0.0_230714.2015_81 2	Internal updates, no customer impact.
Minor	21.11.0.0.0_230714.2015_80 2	Internal updates, no customer impact.
Bundle	21.11.0.0.0_230714.2015_78 6	Bugs fixed: <ul style="list-style-type: none"> • GGS-7261: Diagnostic collection fails to collect when log files are larger than 14GB • Bug 35353317: MySQL - Replicat using incorrect metadata while processing DDL record when source database is case sensitive and target is case insensitive • Bug 32131835: Generic - Enhance logdump to optionally disable all column data output • Bug 34732293: Generic - Mapper memory leak occurs when parallel Replicat is started with ATCSN/AFTERCSN • Bug 35164853: Generic - When the GLOBALS file is not present, the AllowNullableKeys command is effectively true • Bug 35225748: Generic - Extract registration information getting removed from microservices metadata • Bug 35350661: Generic - Modify recvsrvr /targets and /targets/path endpoints • Bug 35401164: Generic - Added a new TRANLOGOPTIONS UNSUPPORTEDLDCROPTION parameter
Minor	21.9.0.0.0_230120.0600_766	Internal updates, no customer impact.
Minor	21.9.0.0.0_230120.0600_765	Bugs fixed: <ul style="list-style-type: none"> • BUG 35256956 - Replicat process not terminating after abend • BUG 35384626 - Extract keeps loading after creating it in the WEBUI
Minor	21.9.0.0.0_230120.0600_722	Internal updates, no customer impact.
Bundle	21.9.0.0.0_230120.0600_716	Bugs fixed: <ul style="list-style-type: none"> • Bug 34821283 - AdminClient running in OCI Cloud Shell fails to connect to MySQL deployment • Bug 34811323 - Extract keeps abending when OGG-00146 (Lost connection to MySQL server during query)
Minor	21.8.0.0.0_221119.1258_703	Internal updates, no customer impact.
Minor	21.8.0.0.0_221119.1258_687	Internal updates, no customer impact.
Bundle	21.8.0.0.0_221119.1258_663	Bugs fixed: <ul style="list-style-type: none"> • Bug 34738804 - Administration Service UI/API hangs intermittently. • Bug 34784214 - Alias not showing under "Credential Alias" when adding a new Extract when username has '@' • Bug 34764925 - Web-UI: Drop-down menu does not expand or show all deployments
Bundle	21.7.0.0.0_220820.1908_654	Internal updates, no customer impact.

Type	Version number	Features and updates included
Minor	21.7.0.0.0_220820.1908_616	Bugs fixed: <ul style="list-style-type: none"> Bug 34382174 - Extract abends upon start on this error, 'ggs::gglib::ggcore::CException'

Oracle GoldenGate for PostgreSQL versions

Learn about what's included in each of the latest Oracle GoldenGate for PostgreSQL versions.

Note:

- To find out when a version was released and how long you have to upgrade, refer to your deployment details page.
- To find out what new features are available, see [What's new in OCI GoldenGate](#).
- For more details on bug fixes and enhancements listed in the table below, see [Oracle GoldenGate bugs fixed and enhancements](#)

Type	Version number	Features and updates included
Bundle	21.12.0.0.0_231125.1006_877	Bugs fixed: <ul style="list-style-type: none"> Bug 35032631: Certify Azure database for PostgreSQL 14: capture and delivery Bug 32408803: Oracle Goldengate for PostgreSQL Change Data Capture (CDC) support for enum datatype Bug 34336451: Set the option keepalive =1 in the connection field Bug 35076537: Extract getting abended with positioning sequence ID out of order error for BYTEA and TEXT data types Bug 35575990: Extract fails to parse the record when colon is part of the column name
Minor	21.11.0.0.0_230722.0716_874	Internal updates, no customer impact.
Minor	21.11.0.0.0_230722.0716_862	Internal updates, no customer impact.
Minor	21.11.0.0.0_230722.0716_841	Internal updates, no customer impact.
Minor	21.11.0.0.0_230722.0716_822	Internal updates, no customer impact.
Minor	21.11.0.0.0_230722.0716_812	Internal updates, no customer impact.
Minor	21.11.0.0.0_230722.0716_802	Internal updates, no customer impact.

Type	Version number	Features and updates included
Bundle	21.11.0.0.0_230722.0716_786	Bugs fixed: <ul style="list-style-type: none"> • GGS-7261: Diagnostic collection failes to collect when log files are larger than 14GB • Bug 35367684: The initial load Extract does not capture the VARCHAR column without any explicit length, defined as unique key • Bug 35615197: The DB login happens successfully with wrong host entry in the userdialias string • Bug 32131835: Generic - Enhance logdump to optionally disable all column data output • Bug 34732293: Generic - Mapper memory leak occurs when parallel Replicat is started with ATCSN/AFTERCSN • Bug 35164853: Generic - When the GLOBALS file is not present, the AllowNullableKeys command is effectively true • Bug 35225748: Generic - Extract registration information getting removed from microservices metadata • Bug 35350661: Generic - Modify recvsrvr /targets and /targets/path endpoints • Bug 35401164: Generic - Added a new TRANLOGOPTIONS UNSUPPORTEDLDCROPTION parameter
Minor	21.10.0.0.0_230518.0147_766	Internal updates, no customer impact.
Minor	21.10.0.0.0_230518.0147_765	Bugs fixed: <ul style="list-style-type: none"> • BUG 35256956 - Replicat process not terminating after abend • BUG 35384626 - Extract keeps loading after creating it in the WEBUI
Bundle	21.10.0.0.0_230518.0147_744	Bugs fixed: <ul style="list-style-type: none"> • Bug 35040942: PostgreSQL - The Extract fails to connect to the Postgresql database when an asterisk exists in the password • Bug 35095323: PostgreSQL - Support multiple hosts and ports for PostgreSQL connections • Bug 35222608: PostgreSQL - Enhance the PostgreSQL WebUI to accept multiple hosts and ports • Bug 35290862: PostgreSQL - Oracle GoldenGate does not detect primary key for tables with mulitple indexes • Bug 35347431: PostgreSQL - Add support for using client certificate and client key for ODBC connection in absence of a keyStore • Bug 34626040: Generic - Replicat abends when trying to use _LOW_WATERMARK_UPDATE_TIME parameter in the parameter file • Bug 35161065: Generic - Modifying distribution path tcpRcvBuf or tcpSndBuf prevents path starting and causes Distsrvr to fail with core dump • Bug 35229639: Generic - ADMINRSVR Web UI: Request ER lag by report lag command • BUG 35384626 - Extract keeps loading after creating it in the WEBUI
Minor	21.9.0.0.0_230228.0458_722	Internal updates, no customer impact.

Type	Version number	Features and updates included
Bundle	21.9.0.0.0_230228.0458_716	Internal updates, no customer impact.
Minor	21.8.0.0.0_221119.1258_703	Internal updates, no customer impact.
Bundle	21.8.0.0.0_221119.1258_687	Internal updates, no customer impact.
Minor	21.7.0.0.0_220731.2140_663	Bugs fixed: <ul style="list-style-type: none"> Bug 34738804 - Administration Service UI/API hangs intermittently. Bug 34764925 - Web-UI: Drop-down menu does not expand or show all deployments

Oracle Data Transforms versions

Learn about what's included in each of the latest Oracle Data Transforms versions.



Note:

- To find out when a version was released and how long you have to upgrade, refer to your deployment details page.
- To find out what new features are available, see [What's new in OCI GoldenGate](#).
- For more details on bug fixes and enhancements listed in the table below, see [What's New in Oracle Data Transforms](#)

Type	Version number	Features and updates included
Minor	23.08.21.0.0_231128.1958_874	See what's included in this release.
Minor	23.06.28.0.0_230822.1141_862	Initial Limited Availability release

Oracle Cloud Infrastructure GoldenGate Events

Events are structured messages that indicate changes in resources. You can create rules to subscribe to events and stay informed of any abnormal activity among your resources, when it's time to upgrade your deployment, or purge unused files.

Deployment Event Types

OCI GoldenGate deployments emit the following event types:

Friendly name	Event type
Create Deployment Begin	<code>com.oraclecloud.goldengate.CreateDeployment.begin</code>

Friendly name	Event type
Create Deployment End	<code>com.oraclecloud.goldengate.CreateDeployment.end</code>
Delete Deployment Begin	<code>com.oraclecloud.goldengate.DeleteDeployment.begin</code>
Delete Deployment End	<code>com.oraclecloud.goldengate.DeleteDeployment.end</code>
Restore Deployment Begin	<code>com.oraclecloud.GoldenGate.RestoreDeployment.begin</code>
Restore Deployment End	<code>com.oraclecloud.GoldenGate.RestoreDeployment.end</code>
Start Deployment Begin	<code>com.oraclecloud.GoldenGate.StartDeployment.begin</code>
Start Deployment End	<code>com.oraclecloud.GoldenGate.StartDeployment.end</code>
Stop Deployment Begin	<code>com.oraclecloud.GoldenGate.StopDeployment.begin</code>
Stop Deployment End	<code>com.oraclecloud.GoldenGate.StopDeployment.end</code>
Update Deployment Begin	<code>com.oraclecloud.GoldenGate.UpdateDeployment.begin</code>
Update Deployment End	<code>com.oraclecloud.GoldenGate.UpdateDeployment.end</code>
Storage utilization	<code>com.oraclecloud.GoldenGate.StorageUtilization</code>

Here's a reference event for deployments:

```
{
  "eventType": "com.oraclecloud.goldengate.updatedeployment.begin",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "GoldenGate",
```

```

    "eventTime": "2021-06-23T12:00:00.000Z",
    "contentType": "application/json",
    "data": {
      "compartmentId": "ocidl.compartment.oc1..unique_ID",
      "compartmentName": "example_compartment",
      "resourceName": "example_name",
      "resourceId":
"ocidl.goldengatedeployment.<realm>.<region>..<unique_ID>",
      "availabilityDomain": "availability_domain",
      "freeFormTags": {},
      "definedTags": {}
    },
    "eventID": "unique_ID",
    "extensions": {
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    }
  }
}

```

Database Registration Event Types

OCI GoldenGate database registrations emit the following event types:

Friendly name	Event type
Create Database Registration Begin	com.oraclecloud.goldengate.CreateDatabaseRegistration.begin
Create Database Registration End	com.oraclecloud.goldengate.CreateDatabaseRegistration.end
Delete Database Registration Begin	com.oraclecloud.GoldenGate.DeleteDatabaseRegistration.begin
Delete Database Registration End	com.oraclecloud.GoldenGate.DeleteDatabaseRegistration.end
Update Database Registration Begin	com.oraclecloud.GoldenGate.UpdateDatabaseRegistration.begin
Update Database Registration End	com.oraclecloud.GoldenGate.UpdateDatabaseRegistration.end

Here's a reference event for database registrations:

```

{
  "eventType":
"com.oraclecloud.goldengate.createdatabaseregistration.begin",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
}

```

```

"source": "GoldenGate",
"eventTime": "2021-06-23T12:00:00.000Z",
"contentType": "application/json",
"data": {
  "compartmentId": "ocidl.compartment.oc1..unique_ID",
  "compartmentName": "example_compartment",
  "resourceName": "example_name",
  "resourceId": "ocidl.goldengatedeployment.<realm>.<region>..<unique_ID>",
  "availabilityDomain": "availability_domain",
  "freeFormTags": {},
  "definedTags": {}
},
"eventID": "unique_ID",
"extensions": {
  "compartmentId": "ocidl.compartment.oc1..unique_ID"
}
}

```

Deployment Backup Event Types

OCI GoldenGate deployment backups emit the following event types:

Friendly name	Event type
Create Deployment Backup Begin	com.oraclecloud.GoldenGate.CreateDeploymentBackup.begin
Create Deployment Backup End	com.oraclecloud.GoldenGate.CreateDeploymentBackup.end
Delete Deployment Backup Begin	com.oraclecloud.GoldenGate.DeleteDeploymentBackup.begin
Delete Deployment Backup End	com.oraclecloud.GoldenGate.DeleteDeploymentBackup.end
Update Deployment Backup Begin	com.oraclecloud.GoldenGate.UpdateDeploymentBackup.begin
Update Deployment Backup End	com.oraclecloud.GoldenGate.UpdateDeploymentBackup.end

Here's a reference event for Deployment Backups:

```

{
  "eventType": "com.oraclecloud.goldengate.deletedeploymentbackup.begin",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "GoldenGate",

```

```

    "eventTime": "2021-06-23T12:00:00.000Z",
    "contentType": "application/json",
    "data": {
      "compartmentId": "ocidl.compartment.oc1..unique_ID",
      "compartmentName": "example_compartment",
      "resourceName": "example_name",
      "resourceId":
    "ocidl.goldengatedeployment.<realm>.<region>..<unique_ID>",
      "availabilityDomain": "availability_domain",
      "freeFormTags": {},
      "definedTags": {}
    },
    "eventID": "unique_ID",
    "extensions": {
      "compartmentId": "ocidl.compartment.oc1..unique_ID"
    }
  }
}

```

Deployment Lifecycle Event Types

OCI GoldenGate deployment lifecycle states emit the following events:

Friendly name	Event type
GoldenGate Deployment Active	com.oraclecloud.goldengate.stateactive
GoldenGate Deployment Creating	com.oraclecloud.goldengate.statecreating
GoldenGate Deployment Deleted	com.oraclecloud.goldengate.statedeleted
GoldenGate Deployment Deleting	com.oraclecloud.goldengate.statedeleting
GoldenGate Deployment Failed	com.oraclecloud.goldengate.statefailed
GoldenGate Deployment Inactive	com.oraclecloud.goldengate.stateinactive
GoldenGate Deployment Needs Attention	com.oraclecloud.goldengate.stateneedsattention
GoldenGate Deployment Updating	com.oraclecloud.goldengate.stateupdating

Here's a reference event for deployment lifecycle state:

```
{
  "eventType": "com.oraclecloud.goldengate.stateactive",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "GoldenGate",
  "eventTime": "2021-06-23T12:00:00.000Z",
  "contentType": "application/json",
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "compartmentName": "example_compartment",
    "resourceName": "example_name",
    "resourceId": "ocidl.goldengatedeployment.<realm>.<region>..<unique_ID>"
  },
  "eventID": "unique_ID",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  }
}
```

Upgrade Notification Event Type

The OCI GoldenGate upgrade notification emits the following event type:

Friendly name	Event type
GoldenGate Upgrade Notification	com.oraclecloud.goldengate.upgradenotification

Here's the reference event for an upgrade notification:

```
{
  "eventType": "com.oraclecloud.goldengate.upgradenotification",
  "cloudEventsVersion": "0.1",
  "eventTypeVersion": "2.0",
  "source": "GoldenGate",
  "eventTime": "2021-06-23T12:00:00.000Z",
  "contentType": "application/json",
  "data": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID",
    "compartmentName": "example_compartment",
    "resourceName": "example_name",
    "resourceId": "ocidl.goldengatedeployment.<realm>.<region>..<unique_ID>",
    "availabilityDomain": "availability_domain",
    "additionalDetails": {
      "message": "notification message"
    }
  },
  "eventID": "unique_ID",
  "extensions": {
    "compartmentId": "ocidl.compartment.oc1..unique_ID"
  }
}
```

```
}  
}
```

Oracle Cloud Infrastructure GoldenGate Metrics

Monitor the deployment health, capacity, and overall performance of your Oracle Cloud Infrastructure GoldenGate deployments using metrics, alarms, and notifications. In the Oracle Cloud console, you can monitor metrics on the OCI GoldenGate deployment details page or using the Metrics Explorer.

Resources: `goldengate-deployments`, `goldengate-connections`

Overview

Oracle Cloud Infrastructure GoldenGate metrics help you measure the amount of data replicated between source and target databases.

The following terms are helpful for understanding metrics:

- **Namespace:** A container for Oracle Cloud Infrastructure GoldenGate metrics. The namespace for Oracle Cloud Infrastructure GoldenGate is `oci_goldengate`.
- **Metrics:** The fundamental concept in telemetry and monitoring. Metrics define a time-series set of datapoints. Each metric is uniquely defined by namespace, metric name, compartment identifier, a set of one or more dimensions, and a unit of measure. Each datapoint has a timestamp, a value, and a count associated with it.
- **Dimensions:** A key-value pair that defines the characteristics associated with the metric. For example, `resourceId`, which is the Oracle Cloud Infrastructure GoldenGate deployment OCID.
- **Statistics:** Metric data aggregations over specified periods of time. Aggregations are done using the namespace, metric name, dimensions, and the datapoint unit of measure within the time period specified.
- **Alarms:** Used to automate operations monitoring and performance. An alarm keeps track of changes that occur over a specific period of time. It also performs one or more defined actions, based on the rules defined for the metric.

Prerequisites

- **IAM policies:** To monitor resources, you must be given the required type of access in a policy written by an administrator, whether you're using the Console or the REST API with an SDK, CLI, or other tool. The policy must give you access to the monitoring services as well as the resources being monitored. If you try to perform an action and get a message that you don't have permission or are unauthorized, confirm with your administrator the type of access you've been granted and which compartment you should work in. For more information on user authorizations for monitoring, see [Monitoring](#) or [Notifications](#).
- The metrics listed on this page are automatically available for any Oracle Cloud Infrastructure GoldenGate deployment you create. You do not need to enable monitoring on the resource to get these metrics.

Available Metrics



Note:

Ensure that you upgrade your deployment to the latest version to leverage all available metrics.

Oracle Cloud Infrastructure GoldenGate metrics may include the following dimensions:

- `deploymentId`: For all metrics, the `deploymentId` is the deployment OCID.
- `deploymentName`: Name of the Oracle Cloud Infrastructure GoldenGate deployment.
- `ExtractName`: Name of an Extract process in the Oracle Cloud Infrastructure GoldenGate deployment.
- `ReplicatName`: Name of a Replicat process in the Oracle Cloud Infrastructure GoldenGate deployment.
- `DistributionPathName`: Name of a Distribution Path process in the Oracle Cloud Infrastructure GoldenGate deployment.
- `ReceiverPathName`: Name of a Receiver Path process in the Oracle Cloud Infrastructure GoldenGate deployment.

Metric	Metric Description	Metric Dimensions	Recommended Actions
Name: <code>CpuUtilization</code> Display Name: CPU Utilization	Total CPU usage percentage by all consumer groups. Check the CPU Utilization when: <ul style="list-style-type: none"> • The OCI GoldenGate deployment console is slow or not responsive • There is an Extract or Replicat lag 	<code>deploymentId</code> <code>deploymentName</code>	If CPU Utilization is high, you can: <ul style="list-style-type: none"> • Add OCPUs to your deployment • Enable Autoscale, if not already enabled • Check storage utilization and trail size, and purge trail files if high
Name: <code>OcpuConsumption</code> Display Name: OCPU Consumption	Total number of OCPUs used by the deployment. When the count is lower than the minimum number of OCPUs, the minimum is shown. When the number of OCPUs is greater than the minimum number, the actual number of OCPUs used is shown.	<code>deploymentId</code> <code>deploymentName</code>	If OCPU Consumption is high, you can: <ul style="list-style-type: none"> • Add OCPUs to your deployment • Enable Autoscale, if not already enabled

Metric	Metric Description	Metric Dimensions	Recommended Actions
Name: MemoryUtilization Display Name: Memory Utilization	<p>Percentage of available memory used.</p> <p>The need for memory is aligned with the size of the data replicated. If enough memory is allocated, then each open transaction is kept in memory until a commit record is received.</p>	deploymentId deploymentName	<p>If Memory Utilization is high, you can:</p> <ul style="list-style-type: none"> • Add OCPUs to your deployment • Enable Autoscale, if not already enabled
Name: DeploymentHealth Display Name: Overall Deployment Health	<p>Overall percentage health of deployment services.</p> <p>There are four services: Administration service, Distribution service, Receiver service, and Performance Metric service. If all four are running healthy, the expected score is 100%. If Deployment Health is 50%, then only two of the services are running healthy.</p>	deploymentId deploymentName	<p>When you create Extract, Replicat, Distribution or Receiver Paths, you can mark the process as Critical to Deployment Health under Managed Options. If the Deployment Health is >100%, then check the processes marked as Critical to Deployment Health.</p>
Name: DeploymentInboundLag Display Name: Deployment Inbound Lag	<p>Average lag, in seconds, for all inbound streams critical to deployment health</p>	deploymentId deploymentName	Not applicable
Name: DeploymentOutboundLag Display Name: Deployment Outbound Lag	<p>Average lag, in seconds, for all outbound streams critical to deployment health</p>	deploymentId deploymentName	Not applicable
Name: SwapSpaceUsage Display Name: Swap Space Usage	<p>Percentage of Swap Space used by the deployment.</p> <p>As OCI GoldenGate only writes only committed transaction to the trail files, all the uncommitted transactions are cached in memory. Cache uses both physical memory and swap space (virtual memory). Swap space is located on hard drives to provide additional memory when the physical memory (RAM) is full.</p>	deploymentId deploymentName	<p>If Swap Space Usage is increasing, consider adding additional OCPUs to the deployment to increase physical memory (RAM).</p>

Metric	Metric Description	Metric Dimensions	Recommended Actions
Name: TempSpaceUsage Display Name: Temporary Space Usage	Percentage of temporary space used by the deployment. When total cached transaction data exceeds the Cachesize setting, Extract writes cache data to temporary files. It is more efficient for the operating system to swap to disk than it is for Extract to write temporary files.	deploymentId deploymentName	If Temp Space Usage is increasing, consider adding additional OCPUs to the deployment to increase physical memory (RAM).
Name: FileSystemUsage Display Name:	Percentage of File System Space used by the deployment	deploymentId deploymentName	If File System Usage is high: <ul style="list-style-type: none"> • Check trail file size and purge unnecessary trail files • Check Temp Space Usage to see if OCI GoldenGate ran short of physical memory (RAM)
Name: ExtractStatus Display Name: Extract Status	Health percentage of an Extract process in the deployment <ul style="list-style-type: none"> • 100% when process is Running • 0% when process is Abended or Stopped 	deploymentId deploymentName ExtractName	If an Extract processes is abended or stopped, check the report file for the root cause or error to troubleshoot the issue.
Name: ReplicatStatus Display Name: Replicat Status	Health percentage of a Replicat process in the deployment <ul style="list-style-type: none"> • 100% when process is Running • 0% when process is Abended or Stopped 	deploymentId deploymentName ReplicatName	If a Replicat is stopped or abended, then check the Replicat report file for the root cause or error to diagnose issues.
Name: DistributionPathStatus Display Name: Distribution Path Status	Health percentage of a Distribution Path process in the deployment <ul style="list-style-type: none"> • 100% when process is Running • 0% when process is Abended or Stopped 	deploymentId deploymentName DistributionPathName	If abended or stopped, then possible causes are: <ul style="list-style-type: none"> • Change in credentials • Receiver service stopped • Target deployment stopped
Name: ReceiverPathStatus Display Name: Receiver Path Status	Health percentage of a Receiver Path process in the deployment <ul style="list-style-type: none"> • 100% when process is Running • 0% when process is Abended or Stopped 	deploymentId deploymentName ReceiverPathName	If the Receiver Path Status is stopped or abended: <ul style="list-style-type: none"> • Check for changes in credentials • Check the target deployment health • Check for network issues between the source and target deployments

Metric	Metric Description	Metric Dimensions	Recommended Actions
Name: ExtractLag Display Name: Extract Lag	The difference, in seconds, between the time the Extract processed a record (based on the system clock) and the timestamp of that record in the data source.	deploymentId deploymentName ExtractName	If the Extract Lag is high, then: <ul style="list-style-type: none"> • Check CPU Utilization to see if the deployment ran out of resources • Check Memory Utilization to see if the assigned resources can handle the Extract size • If the data source is an on-premises database, check network health and latency • Check for performance issues with the source database • Check file system storage
Name: ReplicatLag Display Name: Replicat Lag	The difference, in seconds, between the time the Replicat processed the last record (based on the system clock) and the timestamp of the record in the trail.	deploymentId deploymentName ReplicatName	If the Replicat Lag is high: <ul style="list-style-type: none"> • Check CPU and Memory Utilization to see if the deployment ran out of resources • Check Extract Lag for any latency issues • Check the Distribution Path Lab for latency issues • Check network latency between OCI GoldenGate and the target database. Recommended roundtrip ping is 5ms or less. • HANDLECOLLISIONS can cause performance issues and not recommended for Change Data Capture (CDC) replication
Name: DistributionPathLag Display Name: Distribution Path Lag	Average lag, in seconds, of a Distribution Path process in the deployment. For example, if the source and target deployments are running in two different data centers, network latency issues could impact lag.	deploymentId deploymentName DistributionPathName	Not applicable

Metric	Metric Description	Metric Dimensions	Recommended Actions
Name: ReceiverPathLag	Average lag, in seconds, of Receiver Path process in the deployment	deploymentId deploymentName ReceiverPathName	Not applicable
Display Name: Receiver Path Lag			

Using the Console

To view Oracle Cloud Infrastructure GoldenGate metrics:

1. In the Console navigation menu, under **Solutions and Platform**, go to **Monitoring** and then select **Service Metrics**.
2. For **Compartment**, select the compartment that contains the Oracle Cloud Infrastructure GoldenGate deployments you're interested in.
3. For **Metric Namespace**, select **oci_goldengate**.

Refresh your browser to view the latest metrics emitted by the service.

Oracle Cloud Infrastructure GoldenGate Policies

To control access to Oracle Cloud Infrastructure GoldenGate and the type of access each user group has, you must create policies.

For example, you can create an Administrators group whose members can access all OCI GoldenGate resources. You can then create a separate group for everyone else who's involved with OCI GoldenGate, and create policies that restricts their access to OCI GoldenGate resources in different compartments.

For a complete list of Oracle Cloud Infrastructure policies, see [policy reference](#).

Create policies

Policies define what actions members of a group can perform, and in which compartments.

You create policies using the Oracle Cloud console. In the Oracle Cloud console navigation menu, go to **Identity & Security**, and then under **Identity**, and click **Policies**. Policies are written in the following syntax:

```
allow group <identity-domain>/<group-name> to <verb> <resource-type> in
<location> where <condition>
```

- **<identity-domain>:** (Optional) If using OCI IAM for identity management, then include the identity domain of the user group. If omitted, then OCI uses the default domain.
- **<group-name>:** The name of the user group you're giving permissions to
- **<verb>:** Gives the group a certain level of access to a resource-type. As the verbs go from *inspect* to *read* to *use* to *manage*, the level of access increases and the permissions granted are cumulative.
To learn about the relationship between permissions and verbs, see [Permissions](#).
- **<resource-type>:** The type of resource you're giving a group permission to work with. There are individual resources, such as `goldengate-deployments` and `goldengate-`

connections, and there are resource families, such as `goldengate-family`, which includes both `goldengate-deployments` and `goldengate-connections`. For more information, see [resource-types](#).

- `<location>`: Attaches the policy to a compartment or tenancy. You can specify a single compartment or compartment path by name or OCID, or specify `tenancy` to cover the entire tenancy.
- `<condition>`: Optional. One or more conditions for which this policy will apply.

Learn more about [policy syntax](#).

How to create a policy

To create a policy:

1. In the Console navigation menu, under **Governance and Administration**, go to **Identity**, and then click **Policies**.
2. Click **Create Policy**.
3. Enter a name and description for the policy.
4. In the **Statement** field, enter a policy rule in the following format:

```
allow <subject> to <verb> <resource-type> in <location> where  
<condition>
```

Conditions are optional. See Details for Verbs + Resource-Type Combinations.

5. (Optional) To add another statement, click **+ Another Statement**.
6. Click **Create**.

For more information about policies, see [how policies work](#), [policy syntax](#), and [policy reference](#).

Minimum recommended policies

At minimum, you need policies to:

- Allow users to *use* or *manage* GoldenGate resources, so that they can work with deployments and connections. For example:

```
allow group <identity-domain>/<group-name> to manage goldengate-  
family in compartment <compartment-name>
```

- Allow users to *manage* network resources, so that they can view and select compartments and subnets, and create and delete private endpoints when creating GoldenGate resources. For example:

```
allow group <identity-domain>/<group-name> to manage virtual-  
network-family in compartment <compartment-name>
```

Optionally, you can further secure network resources using a combination of granular policies. See Policy Examples for Securing Network Resources.

- Allow users to read the Identity and Access Management (IAM) user and group for validations in IAM enabled tenancies:

```
allow service goldengate to {idcs_user_viewer, domain_resources_viewer}
in tenancy
```

- Oracle Vault, to access customer managed encryption keys. For example:

```
allow group <identity-domain>/<group-name> to manage secret-family in
<location>
allow group <identity-domain>/<group-name> to use keys in <location>
allow group <identity-domain>/<group-name> to use vaults in <location>
allow service goldengate to use keys in <location>
allow service goldengate to use vaults in <location>
```

Depending on whether you intend to use the following services, you may also need to add policies for:

- Oracle Databases, for your source and/or target databases. For example:

```
allow group <identity-domain>/<group-name> to read database-family in
compartment <compartment-name>
```

```
allow group <identity-domain>/<group-name> to read autonomous-database-
family in compartment <compartment-name>
```

- Oracle Object Storage, to store manual OCI GoldenGate backups. For example:

```
allow group <identity-domain>/<group-name> to manage objects in <location>
allow group <identity-domain>/<group-name> to inspect buckets in
<location>
```

- OCI Logging, to access log groups. For example:

```
allow group <identity-domain>/<group-name> to manage log-groups in
compartment <compartment-name>
allow group <identity-domain>/<group-name> to manage log-content in
compartment <compartment-name>
```

The following statement gives a group permission to manage tag-namespaces and tags for workspaces:

```
allow group <identity-domain>/<group-name> to manage tag-namespaces in
compartment <compartment-name>
```

To add a defined tag, you must have permission to use the tag namespace. To learn more about tagging, see [Resource Tags](#).

For more information and additional example policies, see OCI GoldenGate Policies.

Policy Examples for Securing Network Resources

You can easily allow users access to network resources within a compartment with the policy:

```
allow group <group-name> to use virtual-network-family in compartment
<compartment-name>
```

Alternatively, you can use the following policies to secure network resources at a more granular level:

Operation	Required Access on Underlying Resources
Create a private endpoint	<p>For the private endpoint compartment:</p> <ul style="list-style-type: none"> • Create VNIC (VNIC_CREATE) • Delete VNIC (VNIC_DELETE) • Update members in a network security group (NETWORK_SECURITY_GROUP_UPDATE_MEMBERS) • Associate a network security group (VNIC_ASSOCIATE_NETWORK_SECURITY_GROUP) <p>For the subnet compartment:</p> <ul style="list-style-type: none"> • Attach subnet (SUBNET_ATTACH) • Detach subnet (SUBNET_DETACH)
Update a private endpoint	<p>For the private endpoint compartment:</p> <ul style="list-style-type: none"> • Update VNIC (VNIC_UPDATE) • Update members in a network security group (NETWORK_SECURITY_GROUP_UPDATE_MEMBERS) • Associate a network security group (VNIC_ASSOCIATE_NETWORK_SECURITY_GROUP)
Delete a private endpoint	<p>For the private endpoint compartment:</p> <ul style="list-style-type: none"> • Delete VNIC (VNIC_DELETE) • Update members in a network security group (NETWORK_SECURITY_GROUP_UPDATE_MEMBERS) <p>For the subnet compartment:</p> <ul style="list-style-type: none"> • Detach subnet (SUBNET_DETACH)
Change a private endpoint compartment	If moving from one compartment to another, all permissions in the original compartment must also be present in the new compartment.

Resource-Types

Oracle Cloud Infrastructure GoldenGate offers both aggregate and individual resource-types for writing policies.

Aggregate Resource-Type	Individual Resource-Types
goldengate-family	goldengate-deployments goldengate-deployment-backups goldengate-connections goldengate-connection-assignments

The APIs covered for the aggregate `goldengate-family` resource-type also cover the APIs for each of the individual resource-types. For example,

```
allow group gg-admins to manage goldengate-family in compartment
<compartment-name>
```

is the same as writing the following policies:

```
allow group gg-admins to manage goldengate-deployments in compartment
<compartment-name>
allow group gg-admins to manage goldengate-connections in compartment
<compartment-name>
allow group gg-admins to manage goldengate-connection-assignments in
compartment <compartment-name>
allow group gg-admins to manage goldengate-deployment-backups in compartment
<compartment-name>
```

Supported Variables

When you add conditions to your policies, you can use either Oracle Cloud Infrastructure general or service specific variables.

Oracle Cloud Infrastructure GoldenGate supports all general variables. For more information, see [general variables for all requests](#).

Details for Verbs + Resource-Type Combinations

There are various Oracle Cloud Infrastructure verbs and resource-types that you can use when you create a policy.

The following tables show the [permissions](#) and API operations covered by each verb for Oracle Cloud Infrastructure GoldenGate. The level of access is cumulative as you go from `inspect` to `read` to `use` to `manage`.

goldengate-deployments

Permission	APIs Fully Covered
INSPECT	
GOLDENGATE_DEPLOYMENT_INSPECT	ListDeployments
READ	
INSPECT +	INSPECT+
GOLDENGATE_DEPLOYMENT_READ	GetDeployment

Permission	APIs Fully Covered
USE	
READ +	READ +
GOLDENGATE_DEPLOYMENT_UPDATE	UpdateDeployment StartDeployment StopDeployment RestoreDeployment
MANAGE	
USE +	USE +
GOLDENGATE_DEPLOYMENT_CREATE	CreateDeployment GetWorkRequest ListWorkRequests ListWorkRequestErrors ListWorkRequestLogs
GOLDENGATE_DEPLOYMENT_DELETE	DeleteDeployment
GOLDENGATE_DEPLOYMENT_MOVE	ChangeDeploymentCompartment

goldengate-connections

Permission	APIs Fully Covered
INSPECT	
GOLDENGATE_CONNECTION_INSPECT	ListConnections
READ	
INSPECT +	INSPECT+
GOLDENGATE_CONNECTION_READ	GetConnection
USE	
READ +	READ +
GOLDENGATE_CONNECTION_UPDATE	UpdateConnection
MANAGE	
USE +	USE +
GOLDENGATE_CONNECTION_CREATE	CreateConnection
GOLDENGATE_CONNECTION_DELETE	DeleteConnection
GOLDENGATE_CONNECTION_MOVE	ChangeConnectionCompartment

goldengate-connection-assignments

Permission	APIs Fully Covered
INSPECT	
GOLDENGATE_CONNECTION_ASSIGNMENT_INSPECT	ListConnectionAssignments
READ	
INSPECT +	INSPECT+
GOLDENGATE_CONNECTION_ASSIGNMENT_READ	GetConnectionAssignment

Permission	APIs Fully Covered
USE	
READ +	READ +
n/a	n/a
MANAGE	
USE +	USE +
GOLDENGATE_CONNECTION_ASSIGNMENT_CREATE	CreateConnectionAssignment
GOLDENGATE_CONNECTION_ASSIGNMENT_DELETE	DeleteConnectionAssignment

goldengate-deployment-backups

Permission	APIs Fully Covered
INSPECT	
GOLDENGATE_DEPLOYMENT_BACKUP_INSPECT	ListDeploymentBackups
READ	
INSPECT +	INSPECT+
GOLDENGATE_DEPLOYMENT_BACKUP_READ	GetDeploymentBackup RestoreDeployment
USE	
READ +	READ +
GOLDENGATE_DEPLOYMENT_BACKUP_UPDATE	UpdateDeploymentBackup
MANAGE	
USE +	USE +
GOLDENGATE_DEPLOYMENT_CREATE	CreateDeploymentBackup
GOLDENGATE_DEPLOYMENT_DELETE	DeleteDeploymentBackup
GOLDENGATE_DEPLOYMENT_BACKUP_MOVE	ChangeDeploymentBackupCompartment

Permissions Required for Each API Operation

Here's a list of the API operations for Oracle Cloud Infrastructure GoldenGate in logical order, grouped by resource-type.

The resource-types are `goldengate-deployments`, `goldengate-connections`, and `goldengate-deployment-backups`.

API Operation	Permission
ListDeployments	GOLDENGATE_DEPLOYMENT_INSPECT
CreateDeployment	GOLDENGATE_DEPLOYMENT_CREATE
GetDeployment	GOLDENGATE_DEPLOYMENT_READ
UpdateDeployment	GOLDENGATE_DEPLOYMENT_UPDATE
DeleteDeployment	GOLDENGATE_DEPLOYMENT_DELETE
StartDeployment	GOLDENGATE_DEPLOYMENT_UPDATE

API Operation	Permission
StopDeployment	GOLDENGATE_DEPLOYMENT_UPDATE
RestoreDeployment	GOLDENGATE_DEPLOYMENT_BACKUP_READ and GOLDENGATE_DEPLOYMENT_UPDATE
ChangeDeploymentCompartment	GOLDENGATE_DEPLOYMENT_MOVE
UpgradeDeployment	GOLDENGATE_DEPLOYMENT_UPDATE
ListConnections	GOLDENGATE_CONNECTION_INSPECT
CreateConnection	GOLDENGATE_CONNECTION_CREATE
GetConnection	GOLDENGATE_CONNECTION_READ
UpdateConnection	GOLDENGATE_CONNECTION_UPDATE
DeleteConnection	GOLDENGATE_CONNECTION_DELETE
ChangeConnectionCompartment	GOLDENGATE_CONNECTION_MOVE
ListConnectionAssignments	GOLDENGATE_CONNECTION_ASSIGNMENT_I NSPECT
CreateConnectionAssignment	GOLDENGATE_CONNECTION_ASSIGNMENT_ CREATE, GOLDENGATE_DEPLOYMENT_UPDATE, GOLDENGATE_CONNECTION_UPDATE
GetConnectionAssignment	GOLDENGATE_CONNECTION_ASSIGNMENT_ READ
DeleteConnectionAssignment	GOLDENGATE_CONNECTION_ASSIGNMENT_ DELETE, GOLDENGATE_DEPLOYMENT_UPDATE, GOLDENGATE_CONNECTION_UPDATE
ListDeploymentBackups	GOLDENGATE_DEPLOYMENT_BACKUP_INSPE CT
GetDeploymentBackup	GOLDENGATE_DEPLOYMENT_BACKUP_READ
CreateDeploymentBackup	GOLDENGATE_DEPLOYMENT_BACKUP_CREA TE
UpdateDeploymentBackup	GOLDENGATE_DEPLOYMENT_BACKUP_UPDA TE
CancelDeploymentBackup	GOLDENGATE_DEPLOYMENT_BACKUP_UPDA TE
DeleteDeploymentBackup	GOLDENGATE_DEPLOYMENT_BACKUP_DELE TE
ChangeDeploymentBackupCompartment	GOLDENGATE_DEPLOYMENT_BACKUP_MOVE
GetDeploymentUpgrade	GOLDENGATE_DEPLOYMENT_UPGRADE_REA D
ListDeploymentUpgrades	GOLDENGATE_DEPLOYMENT_UPGRADE_INS PECT
GetWorkRequest	GOLDENGATE_DEPLOYMENT_CREATE
ListWorkRequests	GOLDENGATE_DEPLOYMENT_CREATE
ListWorkRequestErrors	GOLDENGATE_DEPLOYMENT_CREATE
ListWorkRequestLogs	GOLDENGATE_DEPLOYMENT_CREATE

Known Issues in Oracle Cloud Infrastructure GoldenGate

General

Learn about general known issues that apply to the entire service and how to work around them

Oracle GoldenGate REST APIs return 302 redirects to an index page.

You can use the [GoldenGate REST APIs](#) to manage your OCI GoldenGate deployments. For those familiar with Oracle GoldenGate, note that the Service Manager is not exposed in OCI GoldenGate and any calls made to Service Manager won't be able to return.

AdminClient: Unable to negotiate with <ip-address> port 22: no matching host key type found.

When you use AdminClient in Cloud Shell to connect to your deployment, you may encounter the following message:

```
FIPS mode initialized.  
Unable to negotiate with <ip-address> port 22: no matching host key type  
found. Their offer: ssh-ed25519  
Action completed. Waiting until the work request has entered state:  
( 'SUCCEDED', )  
FIPS mode initialized.  
Unable to negotiate with <ip-address> port 22: no matching host key type  
found. Their offer: ssh-ed25519  
Cannot create ssh tunnelnel
```

Workaround: Complete the following steps:

1. Open a new Cloud Shell session.
2. Create a file using the following command:

```
cat .ssh/config
```

3. Enter the following into the `.ssh/config` file, and then save it:

```
HostkeyAlgorithms ssh-rsa,ssh-ed25519  
PubkeyAcceptedKeyTypes ssh-ed25519,ssh-rsa
```

4. If there's an existing `.ssh/known_hosts` file, delete it.
5. Close the Cloud Shell session.
6. Click **Launch Admin Client** on your deployment details page.

Deployment console

Deployment console fails to load

If you enter a fully qualified domain name (FQDN) whose last portion is longer than 11 characters, the deployment console fails to load.

Workaround: Keep the last portion of your FQDN under 11 characters.

The OCI GoldenGate deployment console is not compatible with Safari web browsers.

The Oracle Cloud Infrastructure GoldenGate Deployment Console will not display correctly when accessed using a Safari web browser.

Workaround: Use Chrome or FireFox browsers instead.

Connecting to a credential can take several minutes

In the deployment console's Configuration screen when you attempt to connect to a credential, it can take several minutes for it to connect successfully. Refreshing your screen will only add time to the connection process.

Workaround: This is a known issue that is resolved in GoldenGate build version oggoracle:21.8.0.0.0_221119.1258_663.

Connections

Learn about known issues related to connections and how to work around them.

October release of new Connection types is not yet available in select regions

The latest release of connection types (October 2023) is not yet available in regions where the Stream Analytics deployment type is in Limited Availability. If you don't see any of the following connection types in your region, then switch to a different region, if possible:

- Google BigQuery
- Google Cloud Storage
- Redis
- Amazon Redshift
- Amazon Kinesis
- Elasticsearch
- Google Cloud SQL for SQL Server
- SingleStoreDB
- SingleStoreDB Cloud

Contact your Oracle representative for more information.

Issue with MongoDB Test connection

You may encounter an error when using Test connection with MongoDB connections. You can ignore this error and test MongoDB connections in the OCI GoldenGate deployment console. In the deployment console, open the navigation menu for the Administration Service, click **Configuration**. Your MongoDB connection should be listed as a credential, where you can click **Connect to <alias>** to test the connection.

Create connection for Autonomous Database Dedicated must use Oracle Database type

If creating a connection for an *Autonomous Database Dedicated* instance, then you must select Oracle Database as the type, instead of Oracle Autonomous Database.

To create a connection for Autonomous Database Dedicated:

1. On the Connections page, click **Create connection**.
2. In the Create connections panel, enter a name for the connection, and optionally, a description.
3. Select a compartment in which to create the connection.
4. From the Type dropdown, select **Oracle Database**.
5. Click **Next**.
6. On the Connection details screen, for **Database details**, select **Enter database information**.
7. For **Database connection string**, enter the TCP connection string. You can find the connection string in the tnsnames.ora file downloaded from the wallet package on the Autonomous Database Dedicated details page. Ensure you use the "_low" connection string.
8. Enter the database username and password.
9. **Do not** upload the database wallet.
10. In the Network connectivity section:
 - a. Select **Dedicated endpoint**.
 - b. For **Session mode**, select **Redirect**.
11. Click **Create**.

Action Required for Autonomous Databases that Use mTLS Authentication

When an Autonomous Database wallet is rotated, the OCI GoldenGate connection to this database must be refreshed to retrieve the latest wallet information.

For more information see, [My Oracle Support \(MOS\) Document 2911553.1](#).

To refresh an Autonomous Database connection: Edit and save the connection to the Autonomous Database (Autonomous Transaction Processing or Autonomous Datawarehouse). Saving the connection automatically downloads and refreshes the wallet. No other changes to the connection is needed.

To verify:

1. Launch the deployment console for a deployment that uses the Autonomous Database connection.

2. In the deployment console, open the navigation menu, and then click **Configuration**.
3. On the Credentials screen, observe the Autonomous Database connection string. Before the wallet is refreshed, the connection string looks like the following:

```
ggadmin@(DESCRIPTION=(TRANSPORT_CONNECT_TIMEOUT=3)
(CONNECT_TIMEOUT=60)(RECV_TIMEOUT=120)(retry_count=20)
(retry_delay=3)(address=(protocol=tcps)(port=1522)(host=adb.us-
phoenix-1.oraclecloud.com))
(CONNECT_DATA=(COLOCATION_TAG=ogginstance)
(FAILOVER_MODE=(TYPE=SESSION)(METHOD=BASIC)(OVERRIDE=TRUE))
(service_name=<adb-servicename>_low.adb.oraclecloud.com))
(security=(MY_WALLET_DIRECTORY="/u02/connections/
ocid1.goldengateconnection.oc1.phx.<ocid>/wallet")
(SSL_SERVER_DN_MATCH=TRUE)(ssl_server_cert_dn="CN=adwc.uscom-
east-1.oraclecloud.com,
OU=Oracle BMCS US, O=Oracle Corporation, L=Redwood City,
ST=California,
C=US"))))
```

After the wallet is refreshed, the connection string is updated to look like the following:

```
ggadmin@(DESCRIPTION=(TRANSPORT_CONNECT_TIMEOUT=3)
(CONNECT_TIMEOUT=60)(RECV_TIMEOUT=120)(retry_count=20)
(retry_delay=3)(address=(protocol=tcps)(port=1522)(host=adb.us-
phoenix-1.oraclecloud.com))
(CONNECT_DATA=(COLOCATION_TAG=ogginstance)
(FAILOVER_MODE=(TYPE=SESSION)(METHOD=BASIC)(OVERRIDE=TRUE))
(service_name=<adb-servicename>_low.adb.oraclecloud.com))
(security=(MY_WALLET_DIRECTORY="/u02/connections/
ocid1.goldengateconnection.oc1.phx.<ocid>/wallet")
(SSL_SERVER_DN_MATCH=TRUE)(ssl_server_dn_match=yes)))
```

MySQL database usernames that include an '@' symbol don't appear in Credential Alias list when creating an Extract in the OCI GoldenGate deployment console

For MySQL databases, usernames that include @ symbols are omitted from the Credential Alias list when creating Extracts in the OCI GoldenGate deployment console.

Workaround: Select a different alias from the list and then manually update the Parameter File on the next screen.

Private endpoints not supported for Azure SQL Managed Instance

Although Network connectivity settings appear on the Create connection screens for Azure SQL Managed Instance, private endpoints for Azure SQL Managed Instance are not currently supported.

Workaround: None.

Network timeout affects database connections using private endpoints.

If you're using a private endpoint to connect to a database, then you may encounter network timeouts when starting or stopping Extract processes.

Workaround: You can do one of the following:

- Apply the latest patches from your deployment details page. In the **Deployment Information** section, under **GoldenGate**, for **Version**, click **Upgrade**.
- If you're unable to apply the latest patches at this time, you can update the connection string to include `EXPIRE_TIME=1`. By default, you may have an EZ connection string in Oracle GoldenGate. This connection string needs to be updated in the Oracle GoldenGate Credential to a long connection string as follows:

```
<username>@//<hostname>:1521/<service_name>
<username> @(DESCRIPTION = (EXPIRE_TIME=1) (ADDRESS_LIST = (ADDRESS =
(COMMUNITY = tcp) (PROTOCOL = TCP) (Host = <hostname>) (Port = 1521)))
(CONNECT_DATA = (SERVICE_NAME = <service_name>)))
```

SCAN Proxy doesn't support TLS

While OCI GoldenGate supports Oracle Single Client Access Name (SCAN) hosts and IPs, the SCAN proxy does not support TLS.

Workaround: You can connect to a RAC database using the Database Node IP.

User OCID Mismatch in OCI Object Storage connection (Federated users only)

If a federated user selects Use current user when creating an OCI Object Storage connection, their OCID doesn't match the OCID picked up by the system.

Workaround: When you create an OCI Object Storage connection, ensure that you choose **Specify another user**, and then enter the federated user's OCID.

To find the user OCID, click **Profile** in the Oracle Cloud console global header, and then select the user name. On the User Details page, under User Information, click **Show** for OCID.

GoldenGate processes

Learn about known issues related to GoldenGate processes and how to work around them.

OCI GoldenGate deployment console cannot display custom/non-default named Discard file

Discard files, by default, follow the naming convention `<process-name>.dsc`. You can see all discard files in the OCI GoldenGate deployment console, unless you renamed them. The deployment console doesn't display custom named discard files.

Workaround: Use the Collect diagnostics tool on the deployment details page to access your discard files.

Replicats fail when using Trail file from MongoDB Extract with BINARY_JSON_FORMAT

When a Replicat uses a Trail file generated from a MongoDB Extract with BINARY_JSON_FORMAT in the Extract parameter file, the Replicat fails with the following error:

```
ERROR 2023-08-04 17:13:13.000421 [main] - Unable to decode column 0 :
Input length = 1
    java.nio.charset.MalformedInputException: Input length = 1 at
java.nio.charset.CoderResult.throwException(CoderResult.java:281)
~[?:1.8.0_311]at
java.nio.charset.CharsetDecoder.decode(CharsetDecoder.java:816)
~[?:1.8.0_311] at

oracle.goldengate.datasource.UserExitDataSource.createColumnValue(UserE
xitDataSource.java:1106)
    [ggdbutil-21.9.0.0.3.001.jar:21.9.0.0.3.001] Exception in thread
"main"
    oracle.goldengate.util.GGException: Unable to decode column 0 :
Input length = 1 at

oracle.goldengate.datasource.UserExitDataSource.createColumnValue(UserE
xitDataSource.java:1203)
```

Workaround: When BINARY_JSON_FORMAT is removed from the Extract parameters, the Replicat runs successfully and documents are represented in Extended JSON format.

Remote change data capture Extracts fail for GTID enabled databases

When you create a Change Data Capture Extract process with the Remote option enabled for a MySQL database that uses global transaction identifiers (GTIDs), the Extract process fails and the following error is reported:

```
ERROR   OGG-25192  Trail file '<trail name>' is remote. Only local
trail allowed for this extract.
```

Workaround: On the Parameter file screen of the Change Data Capture Extract, remove the line, TRANLOGOPTIONS ALTLOGDEST REMOTE.

For more information, see [Using Oracle GoldenGate for MySQL](#).

To create Distribution Paths to send data to or pull data from Oracle Cloud Infrastructure GoldenGate, ensure that you add the root certificate to Certificate Management or your client wallet

To send data to or pull data from OCI GoldenGate, you must create a Distribution Server Path or a target initiated path on the Receiver Server in your on-premises or Marketplace Oracle GoldenGate, respectively. You must also add the OCI GoldenGate root certificate or self-signed certificate to your Oracle GoldenGate Certificate Management (Oracle GoldenGate 21c or higher) or client wallet (Oracle GoldenGate 19c). This creates a trusted connection between your Oracle GoldenGate and OCI

GoldenGate deployments. Only WebSocket Secure (WSS) protocol is supported for Distribution and Receiver Server Paths between Oracle GoldenGate and OCI GoldenGate.

A change in the OCI GoldenGate root certificate will cause the Distribution Server Path or a target initiated path on the Receiver Server in your on-premises or Marketplace Oracle GoldenGate to fail and produce the following error:

```
ERROR   OGG-10390  Oracle GoldenGate Receiver Service:  Generic error -1
noticed for endpoint
          wss://<deployment URL>:443/services/v2/sources?trail=<trail name>.
Error description - SSL
          connection unexpectedly closed.
```

Workaround: To fix this issue, update the certificate in the client wallet or Service Manager's Certificate Management screen to use the OCI GoldenGate Deployment Console root certificate.

Learn more:

- For Oracle GoldenGate 19c users, see [Creating a Distribution Server Path User Certificate](#).
- For users of Oracle GoldenGate 21c or higher, see [Create a Trusted Connection Between Oracle GoldenGate and OCI GoldenGate](#).

Only Digest Authentication is currently supported

Oracle Cloud Infrastructure GoldenGate doesn't currently support certificate-based authentication when you use Oracle Cloud Infrastructure GoldenGate as the Distribution Path target.

Workaround: None.