Informatica® Multidomain MDM
9.7.1 Hotfix 7 and later

Zero Downtime Installation Guide for Oracle

Informatica Multidomain MDM Zero Downtime Installation Guide for Oracle
9.7.1 Hotfix 7 and later
March 2019

Publication Date: 2019-12-19

Table of Contents

Preface 4

Informatica Resources 4

- Informatica Network 4
- Informatica Knowledge Base 4
- Informatica Documentation 4
- Informatica Product Availability Matrices 5
- Informatica Velocity 5
- Informatica Marketplace 5
- Informatica Global Customer Support 5

Chapter 1: Configure Zero Downtime 6

Zero Downtime Overview 6

Zero Downtime Replication with Two Systems 7

- Other Replication Scenarios 7

Review the Requirements 8

Port Numbers Used by Oracle GoldenGate 8

Chapter 2: Prepare the Environment for Oracle Databases 9

Ensure the Source and Target Schemas are Identical 9

Edit the Samples 9

Install Zero Downtime 10

- Zero Downtime Scripts and Source Code 10
- Install Informatica MDM Zero Downtime 10

Install Oracle GoldenGate 11

Start the GoldenGate Manager 12

Prepare the Oracle Databases 13

Populate the Schemas 15

Configure and Deploy the Messaging Stream 16

Chapter 3: Troubleshooting 19

Troubleshooting the installation 19

Metadata validation fails 20

Message queue replication is not working 20

Chapter 4: Administrative Tasks 22

Configuring an Additional DBA User in the MDM Hub 22

Resetting the Messaging Stream After an Oracle Flashback 23

Removing ZDT Replication 23

Preface

In a zero downtime (ZDT) environment, you can upgrade Informatica® Multidomain MDM without shutting down the system.

This guide describes how to set up a ZDT environment for Multidomain MDM. For information about upgrading the software, see the *Multidomain MDM Zero Downtime Upgrade Guide* for your database environment.

Informatica Resources

Informatica provides you with a range of product resources through the Informatica Network and other online portals. Use the resources to get the most from your Informatica products and solutions and to learn from other Informatica users and subject matter experts.

Informatica Network

The Informatica Network is the gateway to many resources, including the Informatica Knowledge Base and Informatica Global Customer Support. To enter the Informatica Network, visit https://network.informatica.com.

As an Informatica Network member, you have the following options:

- Search the Knowledge Base for product resources.
- View product availability information.
- Create and review your support cases.
- Find your local Informatica User Group Network and collaborate with your peers.

Informatica Knowledge Base

Use the Informatica Knowledge Base to find product resources such as how-to articles, best practices, video tutorials, and answers to frequently asked questions.

To search the Knowledge Base, visit https://search.informatica.com. If you have questions, comments, or ideas about the Knowledge Base, contact the Informatica Knowledge Base team at KB_Feedback@informatica.com.

Informatica Documentation

Use the Informatica Documentation Portal to explore an extensive library of documentation for current and recent product releases. To explore the Documentation Portal, visit https://docs.informatica.com.

Informatica maintains documentation for many products on the Informatica Knowledge Base in addition to the Documentation Portal. If you cannot find documentation for your product or product version on the Documentation Portal, search the Knowledge Base at https://search.informatica.com.

If you have questions, comments, or ideas about the product documentation, contact the Informatica Documentation team at infa_documentation@informatica.com.

Informatica Product Availability Matrices

Product Availability Matrices (PAMs) indicate the versions of the operating systems, databases, and types of data sources and targets that a product release supports. You can browse the Informatica PAMs at https://network.informatica.com/community/informatica-network/product-availability-matrices.

Informatica Velocity

Informatica Velocity is a collection of tips and best practices developed by Informatica Professional Services and based on real-world experiences from hundreds of data management projects. Informatica Velocity represents the collective knowledge of Informatica consultants who work with organizations around the world to plan, develop, deploy, and maintain successful data management solutions.

You can find Informatica Velocity resources at http://velocity.informatica.com. If you have questions, comments, or ideas about Informatica Velocity, contact Informatica Professional Services at ips@informatica.com.

Informatica Marketplace

The Informatica Marketplace is a forum where you can find solutions that extend and enhance your Informatica implementations. Leverage any of the hundreds of solutions from Informatica developers and partners on the Marketplace to improve your productivity and speed up time to implementation on your projects. You can find the Informatica Marketplace at https://marketplace.informatica.com.

Informatica Global Customer Support

You can contact a Global Support Center by telephone or through the Informatica Network.

To find your local Informatica Global Customer Support telephone number, visit the Informatica website at the following link:
https://www.informatica.com/services-and-training/customer-success-services/contact-us.html.

To find online support resources on the Informatica Network, visit https://network.informatica.com and select the eSupport option.

CHAPTER 1

Configure Zero Downtime

This chapter includes the following topics:

- Zero Downtime Overview, 6
- Zero Downtime Replication with Two Systems, 7
- Review the Requirements, 8
- Port Numbers Used by Oracle GoldenGate, 8

Zero Downtime Overview

When you need to ensure uninterrupted access to master data, implement a zero downtime environment. In a zero downtime environment, you can maintain access to data in the MDM Hub Store while you upgrade Multidomain MDM. You need a source database in a production environment and a target database in a secondary environment. When the data changes in the source database, the changes are replicated to the target database.

When you need to upgrade Multidomain MDM, you make the target database active while you update the source database. After you finish updating Multidomain MDM, you can replicate the changes that occurred in the target database to the source database.

You use Oracle GoldenGate to configure and manage a zero downtime environment for Multidomain MDM. For more information about Oracle GoldenGate, visit the Oracle website.

Zero Downtime Replication with Two Systems

When an organization maintains parallel environments, you run Multidomain MDM on two systems. Data is replicated from the source system to the target system. Multidomain MDM uses Oracle GoldenGate for replication.

The following image shows an example of a source system and a target system:

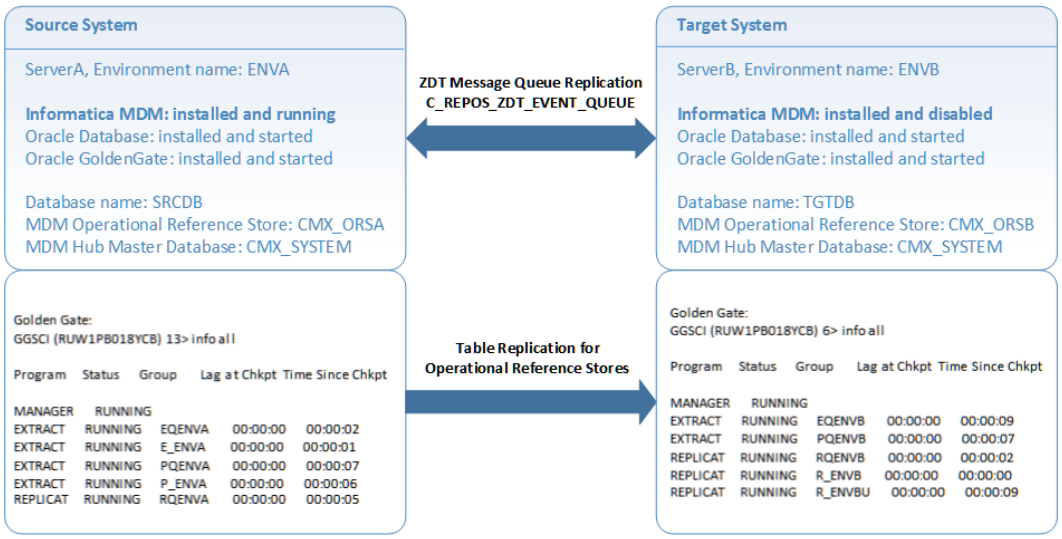

The installed versions of Multidomain MDM and the database software must be the same on the source system and the target system. The name of the MDM Hub Master Database can be different. The Operational Reference Store schema or user name can be different, but the structure of the schemas must be identical.

Oracle GoldenGate replicates the tables in an Operational Reference Store in one direction, from the source system to the target system. The events in the ZDT Message Queue are replicated in both directions.

Other Replication Scenarios

Other types of replications scenarios are possible.

Replication of multiple Operational Reference Store schemas

If you have multiple Operational Reference Store (ORS) schemas within the same Oracle instance, each ORS requires its own set of Oracle GoldenGate processes. Set up the processes for each additional ORS schema in the same way that you set up the first schema.

Replication of schemas that are not ORS schemas

You can replicate other types of schemas by using Oracle GoldenGate. Follow the Oracle GoldenGate documentation to set up replication processes for these schemas.

Replication of non-Informatica tables within the ORS schema

Tables that are not natively part of an ORS are not replicated through Zero Downtime. Follow the Oracle GoldenGate documentation to set up replication processes for these tables.

Review the Requirements

Install the required software on the source system and on the target system. Ensure that you install the same versions of the required software on both systems.

You need the following software:

- Multidomain MDM version 9.7.1 HotFix 7 or later with a supported version of Oracle database
- Microsoft Visual C++ 2005 Redistributable Package (x86) installed as administrator: http://www.microsoft.com/en-us/download/details.aspx?id=3387
- Oracle GoldenGate for Oracle *release*, where *release* matches the installed version of Oracle Database, downloaded from: http://www.oracle.com/technetwork/middleware/goldengate/downloads/index.html

  Important: To verify which versions of Oracle Database and Oracle GoldenGate were supported for your version of Multidomain MDM, see the Product Availability Matrix for the Multidomain MDM version.

  The following table summarizes some of the tested combinations:

Multidomain MDM Version	Oracle GoldenGate Version
9.7.1 HotFix 7	12.1.2.1
10.0 and 10.1	11.2.1.0.18
10.2	12.1.2.1
10.3	12.3.0.1
10.3 HotFix 1	12.3.0.1 or 18.1.0.0.0*

* This version requires that you grant DBA privileges to the GGUser.

Port Numbers Used by Oracle GoldenGate

When the source database and target database reside on different servers, you must open the ports that are used by Oracle GoldenGate. You need one port for Oracle GoldenGate Manager plus one port for each process you run.

For Zero Downtime, you run Oracle GoldenGate Manager and five processes in each environment. Therefore, you need six open ports on the active environment and six open ports on the passive environment. By default, the Oracle GoldenGate port range is 7809-7820.

For information about specifying ports for remote network communications, see the following topics in the *Oracle® GoldenGate Administering Oracle GoldenGate for Windows and UNIX*:

- Maintaining Ports for Remote Connections through Firewalls: https://docs.oracle.com/goldengate/1212/gg-winux/GWUAD/wu_manager.htm#GWUAD142
- Creating the Manager Parameter File: https://docs.oracle.com/goldengate/1212/gg-winux/GWUAD/wu_manager.htm#GWUAD145

CHAPTER 2

Prepare the Environment for Oracle Databases

This chapter includes the following topics:

- Ensure the Source and Target Schemas are Identical, 9
- Edit the Samples, 9
- Install Zero Downtime, 10
- Install Oracle GoldenGate, 11
- Start the GoldenGate Manager, 12
- Prepare the Oracle Databases, 13
- Populate the Schemas, 15
- Configure and Deploy the Messaging Stream, 16

Ensure the Source and Target Schemas are Identical

Back up the Oracle source database that contains the MDM Hub Store. Then copy and restore the backup on the target database. When you are finished, the target database has the same tablespaces as the source database.

Important: During the upgrade, you must drop the source schema, re-create it from the target schema, and then import the database dump file. Do not attempt to bypass this process by applying a change list, because the schemas must be exactly the same in both databases for the replication to work. To avoid making inadvertent changes, enable Production Mode on the source and target databases. Log in to the Hub Console, select the Databases tool, select the database, and enable Production Mode. In future, if you need to apply a change list to the target database, you can disable Production Mode and apply the change list.

Edit the Samples

This guide contains sample code and scripts. To use the samples, copy the samples and substitute your own values. The samples use the forward slash in paths. Ensure that you use the file path separator that works with your operating system.

Install Zero Downtime

Find and run the Zero Downtime scripts on the source system and the target system.

Zero Downtime Scripts and Source Code

You can find the Zero Downtime (ZDT) files in the following directory: `<infamdm>/hub/server/resources/database/oracle`

Use the SQL scripts (.sql files) to update the Operational Reference Store. Use the wrapped source code files (.plb files) to compile the ZDT packages.

The following table describes the purpose of each file:

Script Name	Purpose
compile_types.sql	**Important:** To ensure that the packaging completes successfully, you must run this file before any other file.
cmx_zdt_objects_support.sql	This script performs the following actions: - Creates the C_REPOS_ZDT_EVENT_SEQ repository table. - Creates the sequence object and creates or updates the trigger for the corresponding C_REPAR* table.
cmxlb_pack.plb	Compiles the specification for the `cmxlb` package.
cmx_debug_print_proc.plb	Compiles the DEBUG_PRINT procedure. **Note:** You must run this file before cmxlog_pack.plb.
cmxlog_pack.plb	Compiles the specification for the `cmxlog` package. **Note:** You must run this file before cmxlog_body.plb.
cmxlog_body.plb	Compiles the body for the `cmxlog` package.
cmxzdt_pack.plb	Compiles the specification for the `cmxzdt` package. **Note:** You must run this file before cmxzdt_body.plb.
cmxzdt_body.plb	Compiles the body for the `cmxzdt` package.

Install Informatica MDM Zero Downtime

Use Oracle SQL*Plus (`sqlplus`) to run the SQL scripts and the wrapped source code files. Perform these steps on the source system and the target system.

1. Navigate to the following directory:

   `<MDM installation directory>/hub/server/resources/database/oracle/en_US`

2. Open the `update_javasp.sql` file and verify the following lines of code:

```
HOST loadjava -verbose -force -resolve -oracleresolver -thin -user &ors_name/
&ors_passwd@&tns_name ../siperian-cleansecaller.jar
HOST loadjava -verbose -force -resolve -oracleresolver -thin -user &ors_name/
&ors_passwd@&tns_name ../siperian-dbutil.jar
HOST loadjava -verbose -force -resolve -oracleresolver -thin -user &ors_name/
&ors_passwd@&tns_name ../siperian-zdt.jar
HOST loadjava -verbose -force -resolve -oracleresolver -thin -user &ors_name/
&ors_passwd@&tns_name ../informatica-locale-util.jar
```

If necessary, update the lines of code in the file to match the preceding code.

Note: In MDM 9.7.1 versions, the JAR file `../informatica-locale-util.jar` was named `../informatica-db-locale-util.jar`.

3. Navigate to the following directory:

   `<MDM installation directory>/hub/server/resources/database/oracle/en_US`

4. As the schema user, run the following command on the source database:

   ```
   SQL> @compile_types.sql
   ```

5. Navigate to the following directory:

   `<MDM installation directory>/resources/database/oracle/`

6. Run the files on the source database in the following order:

   ```
   SQL> @cmx_zdt_objects_support.sql
   SQL> @cmxlb_pack.plb
   SQL> @cmx_debug_print_proc.plb
   SQL> @cmxlog_pack.plb
   SQL> @cmxlog_body.plb
   SQL> @cmxzdt_pack.plb
   SQL> @cmxzdt_body.plb
   ```

Install Oracle GoldenGate

Install Oracle GoldenGate on the source system and the target system.

Note: For system requirements and alternative installation instructions, see the Oracle GoldenGate documentation on the Oracle website.

1. Copy the downloaded Oracle GoldenGate .zip file to the system.
2. Extract the .zip file to a local directory named `GGS`.
3. Verify that the following system environment variables are set to the Oracle instance that you use with Informatica MDM: ORACLE_HOME and ORACLE_SID.
4. From a command prompt, navigate to the `GGS` directory.
5. **Windows only**. Enter the following command: `INSTALL ADDSERVICE ADDEVENTS`

   ```
   D:\GGS>INSTALL ADDSERVICE ADDEVENTS
   GoldenGate messages installed successfully.
   Service 'GGSMGR' created.
   Install program terminated normally.
   ```

6. Start the Oracle GoldenGate Command Interpreter:

   ```
   GGS> GGSCI
   ```

7. Create the subdirectories.

   ```
   GGSCI> CREATE SUBDIRS
   ```

The following directories appear under the `GGS` directory:

Directory Name	Contains
`dirprm`	Parameter files
`dirrpt`	Report files
`dirchk`	Checkpoint files
`dirpcs`	Process status files
`dirsql`	SQL script files
`dirdef`	Database definition files
`dirdat`	Extract data files
`dirtmp`	Temporary files
`dirout`	Stdout files

Start the GoldenGate Manager

Start the GoldenGate Manager on the source system and the target system.

1. At a command prompt, create the parameter file.

   ```
   GGSCI> EDIT PARAMS MGR
   ```

2. If you are prompted to create a new file `mgr.prm`, click **Yes**.
3. In the file, add `PORT` *`port_number`*, where *`port_number`* is the port that GoldenGate Manager uses. Save the file.

   Note: The port number that you use must be the same as the value in `C_REPOS_ZDT_STATUS.PUMP_MGRPORT`, which is set in a later step.

4. Start the GoldenGate Manager.

   ```
   GGSCI> START MGR
   ```

5. Verify the port used by the GoldenGate Manager.

   ```
   GGSCI> INFO MGR
   ```

   A message confirms that the GoldenGate Manager is running with the port number that you assigned in the `mgr.prm` file.

   Note: For all subsequent steps, the GoldenGate Manager must be running.

Prepare the Oracle Databases

Prepare the database on the source system and the target system. Before you begin, make sure that you opened the ports required by Oracle GoldenGate, and that the source and target schemas are the same.

Configure the Oracle databases with directory permissions, and then enable replication and logging on the databases. From the Hub Console, enable ZDT on the source schema and the target schema. A DBA must grant execute permissions to the GoldenGate user on both the source system and target system. Do not insert or update data in either database until you finish installing and configuring ZDT successfully.

Tip: If you must update GoldenGate Manager data in the source database while installing ZDT, use system change numbers (SCN) during the data export. Oracle GoldenGate can use the SCNs to replicate the changes after the installation is complete. For more information, see the Oracle GoldenGate documentation.

1. As the SYSTEM user, upgrade the source and target databases by running `update_javasp.sql`.

   The migration process grants privileges automatically for the `GGS_PARAM_DIR` and the `GRANT SELECT ANY DICTIONARY`.

   ```
   SQL> c:/<MDM installation directory>/hub/server/resources/database/oracle/locale/
   update_javasp.sql
   SQL> Create or replace directory GGS_PARAM_DIR as 'D:/GGS/dirprm';
   SQL> Grant read, write on directory GGS_PARAM_DIR to PUBLIC;
   ```

2. As the SYSTEM user, ensure that you have read and write access to the directory `GGS_PARAM_DIR` in the databases.

   ```
   SQL > select * from all_directories where directory_name like 'GGS%';

   OWNER      DIRECTORY_NAME     DIRECTORY_PATH
   ---------- ------------------ ----------------------------------------
   SYS        GGS_PARAM_DIR      D:/ggs/dirprm
   ```

3. **GoldenGate 12 only**. As the SYSDBA user, enable replication on both databases.

   ```
   ALTER SYSTEM SET enable_goldengate_replication=true;
   ```

4. **GoldenGate 12 only**. As the schema user, enable logging for all tables in the schema. Repeat on both databases.

   a. For example, the following script iterates through the tables and creates a `nologging.sql` file in the current directory.

      ```
      spool nologging.sql
      set pagesize 5000
      select 'ALTER TABLE ' || TABLE_NAME || ' LOGGING;' from user_tables where
      LOGGING != 'YES' and TABLE_NAME like 'C^_%' ESCAPE '^';
      exit
      ```

   b. After the script ends, open the `nologging.sql` file, remove headers, and save.

   c. At an sqlplus prompt, log in to the schema as the ORS user and run the edited `nologging.sql` file.

      ```
      SQL> @nologggging.sql
      ```

   d. For each base object in C_REPOS_TABLE, set the `NOLOGGING` indicator to `0`. Repeat on both databases.

      Note: Exclude all tables ending in _STRP. The _STRP tables must have `NOLOGGING` set to `1` for the metadata validation to succeed.

      For example, replace `C_REPOS_TABLE` in the following script with the name of a base object table and run the script:

      ```
      UPDATE C_REPOS_TABLE SET NOLOGGING_IND = 0 WHERE NOLOGGING_IND = 1 and
      table_name != 'C_RBO_BO_CLASS' and table_name != 'C_RBO_BO_HIERARCHY' and
      table_name != 'C_RBO_REL_TYPE';
      ```

5. As the schema user, launch the MDM Hub Console, and enable the ZDT mode for both schemas.

6. Optionally, to verify that the settings are correct, open the repository table `C_REPOS_DB_RELEASE` as the schema user.
   - Ensure that `ZDT_IND` is set to `1`.
   - Ensure that `GLOBAL_NOLOGGING_IND` is set to `0`.
7. As the SYSDBA user, add supplemental log data for the Oracle source database.

```
C:> sqlplus connect sys as sysdba;
SQL> -- make sure you had previously shutdown normally
SQL> STARTUP MOUNT;
SQL> ALTER DATABASE ARCHIVELOG;
SQL> ALTER DATABASE OPEN;
SQL> ALTER DATABASE ADD SUPPLEMENTAL LOG DATA;
SQL> ALTER SYSTEM SWITCH LOGFILE;
SQL> EXIT
```

8. As the schema user, enable the debug log on both the source schema and the target schema.
   a. Start SQL*Plus and log in to the schema.
   b. Open the repository table `C_REPOS_DB_RELEASE`.
   c. Update the columns related to debug and log files with the following values:

Column Name	Value	Description
debug_ind	1	Enables the debug logging process.
debug_level	500	Enables the most detailed level of debugging. When the value is 500, the debug logging process captures errors, warnings, and informational messages, as well as debugging information.
debug_file_name	user-defined	Specifies a name for the file that contains the debug log information.
debug_file_path	user-defined	Specifies the directory where the file is stored. This is the same as the utl_file_dir directory in the database. **Tip:** A DBA creates the utl_file_dir directory within the spfile scope in the Oracle database instance. You can use Oracle Enterprise Manager or SQL*Plus to set the utl_file_dir.
log_file_size	5	Specifies the maximum size of the database log file in MB. The default is 5.
log_file_number	5	Specifies the number of log files used for log rolling. The default is 5.

   d. Repeat on the other schema.

9. **GoldenGate 12 only**. As the SYSDBA user, grant execute permissions to the GoldenGate user on both the source system and target system.

```
SQL> grant Execute on DBMS_STREAMS to PUBLIC;
SQL> grant Execute on DBMS_STREAMS_ADM to PUBLIC;
```

Note: To use Oracle GoldenGate for Oracle version 11.0.2.4 or later, you must grant these execute permissions.

Populate the Schemas

Populate the C_REPOS_ZDT_STATUS repository table on the source system and the target system. Perform all steps as the schema user.

When you set the Local/Remote trail path, ensure that the physical directory exists, otherwise the process will abend. For example, create the `ENVA` and `ENVB` directories in `D:/GGS/dirdat`, and ensure that you have write privileges to those directories.

1. As the schema user, log in to the source schema.
2. Run the following command, substituting appropriate values for the environment name, schema names, directory names, user name, and port:

```
Insert into C_REPOS_ZDT_STATUS
(REPLICATION_TARGET_IND, ACTIVE_UPGRADE_IND, CREATOR, CREATE_DATE
, GGS_HOME_PATH, REPLICAT_NUMBER, TRAIL_FILE_SIZE, DISCARD_FILE_SIZE
, ACCEPTABLE_LAG_MINUTES
, LAG_DETECTION_TOKEN
, LOCAL_ENVIRONMENT_NAME, LOCAL_SCHEMA_NAME, LOCAL_TRAIL_PATH, PUMP_RMTHOST,
PUMP_MGRPORT
, REMOTE_TRAIL_PATH, REMOTE_ENVIRONMENT_NAME, REMOTE_SCHEMA_NAME
, REGULAR_STREAM_ID, EVENT_QUEUE_ID, EXTRACT_PREFIX, REPLICAT_PREFIX
, BATCH_DISABLED_IND, DEFAULT_TIMEOUT_MINUTES, ROWID_ZDT_STATUS,
WRITE_SIF_API_DISABLED_IND)
Values (0, 0, 'MRM Installer', sysdate
, 'D:/GGS', 1, 10, 10, 5
, 'LAG_TOKEN_SRC'
, 'ENVA', 'cmx_ors_a', 'D:/GGS/dirdat/enva/', 'envb_hostname', 9999, 'D:/GGS/dirdat/
envb/', 'ENVB', 'cmx_ors_b'
, 'C', 'Q', 'E','R'
, 0, 40, '1', 0
);
COMMIT;
```

3. Log in to the target schema.
4. Run the following command, substituting appropriate values for the environment name, schema names, directory names, user name, and port:

```
Insert into C_REPOS_ZDT_STATUS
(REPLICATION_TARGET_IND, ACTIVE_UPGRADE_IND, CREATOR, CREATE_DATE
, GGS_HOME_PATH, REPLICAT_NUMBER, TRAIL_FILE_SIZE, DISCARD_FILE_SIZE
, ACCEPTABLE_LAG_MINUTES
, LAG_DETECTION_TOKEN
, LOCAL_ENVIRONMENT_NAME, LOCAL_SCHEMA_NAME, LOCAL_TRAIL_PATH, PUMP_RMTHOST,
PUMP_MGRPORT
, REMOTE_TRAIL_PATH, REMOTE_ENVIRONMENT_NAME, REMOTE_SCHEMA_NAME
, REGULAR_STREAM_ID, EVENT_QUEUE_ID, EXTRACT_PREFIX, REPLICAT_PREFIX
, BATCH_DISABLED_IND
, DEFAULT_TIMEOUT_MINUTES
, ROWID_ZDT_STATUS
, WRITE_SIF_API_DISABLED_IND)
Values
(1, 0, 'MRM Installer', sysdate
, 'D:/GGS', 1, 10, 10, 5
, 'LAG_TOKEN_TGT'
, 'ENVB', 'cmx_ors_b', 'D:/GGS/dirdat/envb/', 'enva_hostname', 9999, 'D:/GGS/dirdat/
enva/', 'ENVA', 'cmx_ors_a'
, 'C', 'Q', 'E','R'
, 1, 40, '1', 1
);
COMMIT;
```

The following table summarizes the status of C_REPOS_ZDT_STATUS repository table:

C_REPOS_ZDT_STATUS Column Name	Source Schema Value	Target Schema Value
REPLICATION TARGET IND	0	1
LOCAL ENVIRONMENT NAME	ENVA	ENVB
LOCAL SCHEMA NAME	cmx_ors_a	cmx_ors_b
LOCAL TRAIL PATH	D:/ggs/dirdat/ENVA/	D:/ggs/dirdat/ENVB/
PUMP RMTHOST	[target host name]	[source host name]
PUMP MGRPORT	[target goldengate mgr port, such as 9999]	[source goldengate mgr port, such as 9999
REMOTE TRAIL PATH	D:/ggs/dirdat/ENVB	D:/ggs/dirdat/ENVA
REMOTE ENVIRONMENT NAME	ENVB	ENVA
REMOTE SCHEMA NAME	cmx_ors_b	cmx_ors_a
REGULAR_STREAM_ID	C	C
EVENT_QUEUE_ID	Q	Q
EXTRACT_PREFIX	E	E
REPLICAT_PREFIX	R	R

Note: The default values from the inserts for the ENVA and ENVB directory for C_REPOS_ZDT_STATUS should be sufficient for most environments.

Configure and Deploy the Messaging Stream

Configure and deploy the messaging stream on both systems.

1. Log in to both the source schema and the target schema.
2. In both schemas, configure the messaging stream:

```
DECLARE
out_error_msg VARCHAR2(32000);
out_return_code INT;
BEGIN
  CMXZDT.CONFIGURE_GGS_EVENT_REPLICAT(CMXZDT.zdt_local);
  CMXZDT.CONFIGURE_GGS_EVENT_EXTRACT(CMXZDT.zdt_local);
  CMXZDT.START_EVENT_QUEUE(out_error_message  => out_error_msg,
                           out_return_code => out_return_code);
  dbms_output.put_line('out_return_code => '||out_return_code||' out_error_msg => '||
substr(out_error_msg,1,250));
END;
/
```

3. On both systems, run the `info all` command from the ggsci prompt. Verify that all event queue processes are running.
4. In the source schema, generate ggsci parameter files and GoldenGate processes:

```
BEGIN
  CMXZDT.CONFIGURE_GGS_EXTRACT(CMXZDT.zdt_local);
END;
 /
```

5. In the target schema, generate ggsci parameter files and GoldenGate processes:

```
BEGIN
  CMXZDT.CONFIGURE_GGS_REPLICAT(CMXZDT.zdt_local);
END;
 /
```

6. In both schemas, set sequences so that they do not overlap:

```
exec cmxzdt.configure_sequences;
```

   This command sets the source schema to even and the target schema to odd to avoid overlaps.
7. In the source schema, start the PUMP and EXTRACT processes:

```
DECLARE
out_error_msg VARCHAR2(32000);
out_return_code INT;
begin
  cmxzdt.start_replication_extract(in_env_type => cmxzdt.zdt_local,
                                   out_error_message  => out_error_msg,
                                   out_return_code => out_return_code);

  dbms_output.put_line('out_return_code => '||out_return_code||' out_error_msg => '||
substr(out_error_msg,1,250));
end;
/
```

8. In the target schema, start the REPLICATE processes:

```
DECLARE
out_error_msg VARCHAR2(32000);
out_return_code INT;
begin
  cmxzdt.start_replication_replay(in_env_type => cmxzdt.zdt_local,
                                   out_error_message => out_error_msg,
                                   out_return_code => out_return_code);

  dbms_output.put_line('out_return_code => '||out_return_code||' out_error_msg => '||
substr(out_error_msg,1,250));
end;
/
```

9. On both systems, run the `info all` command from the ggsci prompt.
10. Verify that the `info all` summary displays ten processes, each with a status of RUNNING.
    - Messaging stream processes have the prefixes EQ, PQ, and RQ.
    - Data extract processes have the prefixes E_ and P_. These processes extract data from the source system.
    - Data replication processes have the prefix R_. These processes replicate data on the target system.

    The following output shows the processes, with one replication process not running:

```
GGSCI (hostname) 13> info all

Program     Status      Group       Lag           Time Since Chkpt

MANAGER     RUNNING
EXTRACT     RUNNING     EQENVA       00:00:00      00:00:00
EXTRACT     RUNNING     EQENVB       00:00:00      00:00:00
EXTRACT     RUNNING     E_ENVA       00:00:00      00:00:08
EXTRACT     RUNNING     PQENVA       00:00:00      00:00:09
```

```
EXTRACT     RUNNING     PQENVB       00:00:00      00:00:00
EXTRACT     RUNNING     P_ENVA       00:00:00      00:00:01
REPLICAT    RUNNING     RQENVA       00:00:00      00:00:00
REPLICAT    RUNNING     RQENVB       00:00:00      00:00:08
REPLICAT    ABENDED     R_ENVB       169:25:21     00:00:02
REPLICAT    RUNNING     R_ENVBU      00:00:00      00:00:07
```

If a process is not running, see the log files in the `D:/ggs/dirrpt` directory. There is a separate .dsc and .rpt file for each process. For more information about troubleshooting process errors, see the Oracle GoldenGate documentation for administrators.

CHAPTER 3

Troubleshooting

This chapter includes the following topics:

- Troubleshooting the installation, 19
- Metadata validation fails, 20
- Message queue replication is not working, 20

Troubleshooting the installation

If you encounter issues while installing Zero Downtime (ZDT), clean your environment and try again. If the problem persists, contact Informatica Global Customer Support.

1. Clean the passive environment.
   a. Remove all the files related to the environment.

```
<GoldenGate install directory>/dirchk
del *

<GoldenGate install directory>/dirdat
del enva/*
del envb/*

<GoldenGate install directory>/dirprm
del eqenv*.prm*
del pqenv*.prm*
del rqenv*.prm*
del e_env*.prm*
del p_env*.prm*
del *.def

<GoldenGate install directory>/dirrpt
del *.dsc
del e_env*.rpt
del eqenv*.rpt
del p_env*.rpt
del pqenv*.rpt
del r_env*.rpt
del rqenv*.rpt
```

   b. Remove the ZDT event queue in the environment.

      Important: Do not drop the GGS_CHECKPOINT and GGS_CHECKPOINT_LOX tables. If you drop these tables, the REPLICAT service does not start.

```
/* Repository tables for ZDT */
delete from C_REPOS_ZDT_EVENT_QUEUE;
```

```
delete from C_REPOS_ZDT_REPLICAT_EXCEPTION;
update C_REPOS_ZDT_ENV_STATE set state = NULL, state_ts = NULL, state_desc =
NULL, updated_by=NULL, update_date=NULL;

/* Tables for Oracle GoldenGate */
delete from GGS_CHECKPOINT;        -- Do not drop this table
delete from GGS_CHECKPOINT_LOX;    -- Do not drop this table

drop table GGS_EVENT_CHECKPOINT cascade constraints;
drop table GGS_EVENT_CHECKPOINT_LOX cascade  constraints;
drop table GGS_EVENT_TRACE cascade  constraints;
commit;
```

c. Search for %GGS% tables. Verify that the GGS_CHECKPOINT and GGS_CHECKPOINT_LOX tables exist. If you find any other tables with GGS in the table name, including TMP_GGS_* tables, you can drop these other tables.

2. Install ZDT.

Metadata validation fails

If the metadata validation fails, check all the tables ending in _STRP. If `NOLOGGING` is set to `1`, change it to `0`.

Message queue replication is not working

If the cmxzdt.configure_ggs_replicate script does not complete, the ZDT message queue replication might not work between the source database and the target database.

1. Check that all the Oracle GoldenGate processes are running. Restart any processes that are not in the RUNNING state.
   In this example, ENVA contains the source database and ENVB contains the target database.

```
EXTRACT RUNNING EQENVA
EXTRACT RUNNING E_ENVA
EXTRACT RUNNING PQENVA
EXTRACT RUNNING P_ENVA
REPLICAT RUNNING RQENVA

EXTRACT RUNNING EQENVB
EXTRACT RUNNING PQENVB
REPLICAT RUNNING RQENVB
REPLICAT ABENDED R_ENVB
REPLICAT RUNNING R_ENVBU
```

   In this example, the R_ENVB process is in the ABENDED state. Try restarting the process.

2. Verify that the message queue replication is working in both directions.
   Insert an event directly into the C_REPOS_ZDT_EVENT_QUEUE table in the source database. Open the same table in the target database. If the event appears in the target database table, replication is working in this direction. Repeat the verification process from the target database to ensure that the replication works in the other direction as well.

   For example, the following code adds an event to the table on ENVA:

```
insert into C_REPOS_ZDT_EVENT_QUEUE ( 'enva', -1, 'test', '', 'envb', 'test',
CURRENT_TIMESTAMP, 'EVENT_TOKEN' );
```

The following code adds an event to the table on ENVB:

```
insert into C_REPOS_ZDT_EVENT_QUEUE ( 'envb', -1, 'test', '', 'enva', 'test',
CURRENT_TIMESTAMP, 'EVENT_TOKEN' );
```

3.

4. If the Oracle GoldenGate processes are running without errors, but the message queue replication is not working, you need to troubleshoot your environment. Navigate to the Oracle GoldenGate directory `dirrpt` and check the .rpt files for information about potential problems.
   For more information about replication issues, see the following Oracle articles on Metalink:
   1. Main Note - Oracle GoldenGate - Troubleshooting (Doc ID 1306476.1)
   2. Master Note - Oracle GoldenGate: Initial Load Techniques and References (Doc ID 1311707.1)
   3. DB Transactions Missing from Oracle GoldenGate Trail Files (Doc ID 1364852.1)
   4. POC for golden gate

CHAPTER 4

Administrative Tasks

This chapter includes the following topics:

- Configuring an Additional DBA User in the MDM Hub, 22
- Resetting the Messaging Stream After an Oracle Flashback, 23
- Removing ZDT Replication, 23

Configuring an Additional DBA User in the MDM Hub

In addition to the owner of the Operational Reference Store and the proxy user, you can configure the MDM Hub so that Oracle GoldenGate uses a different DBA user to connect to the Oracle database. For example, when the connection to the Oracle database is configured with ASM, you need an additional DBA user.

Create the user name and an encrypted password. Use sqlplus to insert the user credentials into the MDM Hub repository table C_REPOS_DB_RELEASE. Add the user name in the db_replication_username column, and add the password in the db_replication_password column. When the columns contain values, Oracle GoldenGate uses these credentials to establish a connection to the database.

To populate the encrypted password, generate an encrypted password using the GGSCI command:

```
GGSCI (in174283) 1> encrypt password <password>
No key specified, using default key...

Encrypted password:  AACAAAAAAAAAAAGAMGMBUEODLIVEUFMJ
```

To generate an encrypted password for C_REPOS_DB_RELEASE.DB_REPLICATION_PASSWORD, use the default key.

For example:

```
GGSCI (in174283) 30> encrypt password password
No key specified, using default key...
Encrypted password:  AACAAAAAAAAAAAIAADHGBAYFMJHITAZF

Plaintext Password: password
Encrypted Password: AACAAAAAAAAAAAIAADHGBAYFMJHITAZF
```

Resetting the Messaging Stream After an Oracle Flashback

After an Oracle flashback, you can reset the messaging stream. Use the `reset_message_queue_trail_files.sql` script, which is included in the replication package.

The script runs the following commands:

```
stop EQENVA
stop PQENVA
DELETE EXTTRAIL D:\ggs\dirdat\enva\qe
ADD EXTTRAIL D:\ggs\dirdat\enva\qe, EXTRACT EQENVA, MEGABYTES 5
DELETE RMTTRAIL D:\ggs\dirdat\envb\qr
ADD RMTTRAIL D:\ggs\dirdat\envb\qr, EXTRACT PQENVA, MEGABYTES 5
start EQENVA
start PQENVA
SEND EXTRACT EQENVA ROLLOVER
SEND EXTRACT PQENVA ROLLOVER
```

The DELETE and ADD commands reset the transaction pointers for the files. The commands do not delete the files.

The ROLLOVER command forces the extract to start a new trail file. For example, if the current file is qe000023, it starts file qe000024.

After the script is finished, start the Oracle GoldenGate Command Interpreter and run the following commands:

```
GGSCI > DBLOGIN USERID [orsid] PASSWORD "[password]"
GGSCI > CLEANUP CHECKPOINTTABLE [orsid]. GGS_EVENT_CHECKPOINT
```

Removing ZDT Replication

To remove the ZDT replication, use the following procedure:

1. Clean up the event queue in Oracle on the source and target schemas:

   ```
   SQL> delete from c_repos_zdt_event_queue;
   Commit;
   ```

2. Connect to GGSCI and issue the following commands:

   ```
   DBLOGIN USERID [orsuser | proxyuser] PASSWORD "[password]"
   STOP *
   DELETE * and answer Y when prompted
   STOP MGR
   ```

3. Connect to ORS user and run the following command:

   ```
   BEGIN
       FOR s IN ( SELECT *
                  FROM   user_log_groups ) LOOP
           EXECUTE IMMEDIATE 'alter table ' || s.table_name
                     || ' drop SUPPLEMENTAL LOG GROUP "' || s.log_group_name || '"';
       END LOOP;
   END;
   /
   ```

   Note: This is required only on the source schema. However, you can execute it on the target as well.

4. Connect to SYSTEM and issue the following command to globally disable supplemental logging:

   ```
   ALTER DATABASE DROP SUPPLEMENTAL LOG DATA;
   ```