Informatica® Multidomain MDM
10.4 HotFix 1

Zero Downtime Upgrade Guide for Oracle

Informatica Multidomain MDM Zero Downtime Upgrade Guide for Oracle
10.4 HotFix 1
September 2020

Publication Date: 2021-10-28

Table of Contents

Preface

Follow the instructions in the Informatica® *Multidomain MDM Zero Downtime Upgrade Guide* to upgrade Multidomain MDM in a zero downtime environment. Zero Downtime is an optionally licensed feature that enables you to minimize disruptions while you upgrade Multidomain MDM.

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.

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

Introduction to Zero Downtime

This chapter includes the following topics:

Zero Downtime Overview

Use Zero Downtime (ZDT) to upgrade Multidomain MDM while providing uninterrupted access to the MDM Hub. Batch and Services Integration Framework (SIF) user processes can run during the ZDT upgrade process.

To use ZDT, you have two environments: a passive environment and an active environment. You use Oracle GoldenGate to manage the data replication and message streams between the environments. While you upgrade the passive environment, MDM users can access the MDM Hub in the active environment. Any change made to the MDM Hub during the upgrade does not impact the ability to replicate changes from the active environment. Backfill mechanisms handle any impact to the MDM Hub metadata caused by the upgrade.

Oracle GoldenGate Processes for ZDT

ZDT requires a set of Oracle GoldenGate processes. Some of the processes extract and replicate data from the active environment to the passive environment.

The following table identifies the prefixes that are used in the process group names:

Prefix	Process Type	Purpose
E_, P_	EXTRACT	Extracts data from the MDM Hub that is located in the active environment.
R_	REPLICAT	Replicates data to the MDM Hub that is located in the passive environment.

For example, the following list shows the processes for ENVA and ENVB, where ENVA is the active environment.

```
GGSCI (hostname) 13> info all

Program     Status      Group
MANAGER     RUNNING
EXTRACT     RUNNING     E_ENVA
REPLICAT    RUNNING     R_ENVB
REPLICAT    RUNNING     R_ENVBU
```

Prerequisites for Zero Downtime

You must identity the source system and target system, ensure the database software is on both systems, install Oracle GoldenGate on both systems, and configure the MDM Hub Store databases for replication. For information about installing ZDT, see the *Multidomain MDM Zero Downtime Installation Guide for Oracle*.

Upgrade with Zero Downtime Workflow

When you upgrade with zero downtime, you perform the following general activities:

1. Upgrade the passive environment.
2. Switch the passive environment and the active environment.
3. Drop the former active environment.
4. Create the new passive environment from a copy of the new active environment.
5. Replicate all data changes that occur during the upgrade so that the environments are the same at the end of the upgrade process.

In this guide, some steps are conditional. Perform these steps only when you are doing a particular type of update or upgrade.

Conditional steps have one or more of the following prefixes:

- **Schema Update**. You changed an Operational Reference Store schema based on business requirements.
- **Schema Update with Data Change**. You changed an Operational Reference Store schema and some data in the database.
- **MDM Upgrade**. You are upgrading Multidomain MDM to a new major, minor, or hotfix release, or you are applying an emergency bug fix.
- **Infrastructure Upgrade**. You are upgrading other software or hardware in the same environments in which Multidomain MDM runs.

You can run the steps for the ZDT upgrade from a command line interface, such as shell scripts, or a command-line job scheduler. The upgrade runs from a single flow of control to allow for almost complete automation of the upgrade process. The ZDT upgrade procedure includes steps for messaging between the active environment and passive environment, maintaining replication control, and integrating backfill.

Upgrading Multidomain MDM from version 9.5.0 or earlier

In Multidomain MDM version 9.5.1, the data structure for the master database schema changed. If you are upgrading from version 9.5.0 or earlier, you must go through a readiness cycle before you start the upgrade

cycle. The readiness cycle identifies data issues that you must resolve before you upgrade. After the readiness cycle, when you upgrade with ZDT, the upgrade cycle updates the data structure in the schema.

To request a version of this guide that contains the steps to upgrade from version 9.5.0 or earlier, contact Informatica Global Customer Support.

CHAPTER 2

Upgrading with Zero Downtime

This chapter includes the following topics:

Transfer the Previous Load Table Data

You can copy the data from the previous load table in the active environment to the previous load table in the passive environment. The data in the previous load table is not transferred during the replication process. If you do not replicate the previous load table, the first stage batch job that you run after the passive environment becomes active might process all the data in the landing table. The names of the previous load tables end in _PRL.

Note: Before you run a stage batch job, ensure that the upgrade is complete. If you run a stage job before the upgrade is complete, the data that the stage batch job adds to the source previous load tables is not added to the target previous load tables.

1. Generate the `prl_expdp.prm` and `prl_impdp.prm` parameter files for the previous load table in the active environment.

   From SQL*Plus in the active environment, run the following command:

   ```
   prl_expdp.prm:

   directory=OUTPUT_DIRECTORY
   logfile=PRL_TABLES_EXPDP_LOG.log
   dumpfile=PRL.dmp
   include=table:"LIKE ''%PRL''"
   CONTENT=DATA_ONLY

   prl_impdp.prm:

   directory=INPUT_DIRECTORY
   logfile=PRL_TABLES_IMPDP_LOG.log
   dumpfile=PRL.dmp
   TABLE_EXISTS_ACTION=APPEND
   CONTENT=DATA_ONLY
   ```

   The parameter files are generated in the `GGS/dirprm` directory.

2. Open a command prompt and navigate to the `GGS/dirprm` directory.

3. Export the data in the previous load table in the active environment:

   ```
   expdp <active Operational Reference Store name>/<TNS password>@<TNS name>
   parfile=prl_expdp.prm
   ```

The `prl.dmp` file is generated in `GGS/dirprm`.

4. Copy the `PRL.dmp` and `prl_impdp.prm` files from the `GGS/dirprm` directory in the active environment to the `GGS/dirprm` directory in the passive environment.
5. If the active and passive environments are not swapping for the first time, truncate the previous load table in the passive environment. You must truncate the previous load table in the passive environment so that the data from the active previous load table can be imported.

   In the passive environment, log in to SQL*Plus and run the following command:

   ```
   TRUNCATE TABLE <Previous Load Table Name>;
   ```
6. In the passive environment, open a command prompt and navigate to `GGS/dirprm`.
7. Run the following command to import the data into the previous load table in the passive environment:

   ```
   impdp <passive Operational Reference Store name>/<TNS password>@<TNS name>
   remap_schema=<active Operational Reference Store name>:<passive Operational
   Reference Store name> parfile=prl_impdp.prm
   ```

Upgrade Steps Controlled from the Passive Environment

To ensure that the systems are synchronized, you must have replication running from the active environment to the passive environment.

Important: Run all the commands in order. Unless otherwise specified, ensure that each process is finished before you run the next command.

1. **MDM Upgrade**. Stop the application server.
2. Disable read services to the passive environment.
3. In the command prompt, to start the upgrade process, enter the following swing command:

   ```
   sip_ant swing_start
   ```

   The swing command sets the `ACTIVE_UPGRADE_IND` to `1`, turns off replication replay, and removes all obsolete objects in the schema of the passive environment.
4. **MDM Upgrade**. Upgrade Multidomain MDM:
   1. Upgrade the Hub Store.
   2. Upgrade the Hub Server.
   3. Upgrade the Process Server.

   For more information about upgrading Multidomain MDM, see the *Multidomain MDM Upgrade Guide*.

   Important: After you finish the upgrade, ensure that the application server runs.
5. **Schema Update**. Run the Metadata Manager command line utility to apply the change list.

   The utility applies the changes in the change list file to the schema. For example, you can use a change list to add or remove columns in the base object table or to set the trust values on columns.
6. Run the Generate Match Tokens batch job to update all the match tokens in the tables that use _STRP as suffix.

You can run the Generate Match Tokens batch job from the Hub Console or by using a service integration framework (SIF) API.

Run From	Steps
Hub Console	1. In the MDM Hub Console, open the Batch Viewer tool. 2. From the Batch Viewer navigation pane, expand the base object for which you want to regenerate all the match tokens. 3. Expand **Generate Match Tokens**. 4. Select the batch job that you want to use to generate match tokens. 5. Select **Re-generate All Match Tokens**. 6. Click **Execute Batch**.
API	To run the Generate Match Tokens batch job on all records, use the ExecuteBatchGenerateMatchTokens request with the `fullRestripInd` attribute set to `1`. The following code sample shows an ExecuteBatchGenerateMatchTokens request to create match tokens for all records in the C_PARTY base object: `<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/" xmlns:urn="urn:siperian.api"> <soapenv:Header/> <soapenv:Body> <urn:executeBatchGenerateMatchTokens> <urn:username>admin</urn:username> <urn:password> <urn:password>admin</urn:password> <urn:encrypted>false</urn:encrypted> </urn:password> <urn:orsId>localhost-orclsnl-UTSOURCE</urn:orsId> <urn:asynchronousOptions> <urn:isAsynchronous>false</urn:isAsynchronous> <urn:jmsReplyTo></urn:jmsReplyTo> <urn:jmsCorrelationId></urn:jmsCorrelationId> </urn:asynchronousOptions> <urn:tableName>C_PARTY</urn:tableName> <urn:fullRestripInd>1</urn:fullRestripInd> </urn:executeBatchGenerateMatchTokens> </soapenv:Body> </soapenv:Envelope>`

7. **MDM Upgrade** and **Schema Update**. Populate the backfill table C_REPOS_ZDT_BACKFILL_TASK to indicate that trust backfill is required on base object tables.

   Use the `sip_ant add_backfill_task` method:

   ```
   sip_ant -Dnoprompt=true -noinput  add_backfill_task -DprocedureName=<backfill type> -DtableName=<base object name> -DusageType=<api> -Dsequence=1
   ```

   where:

   - *backfill type* is the type of backfill. Use one of the following backfill types:
     - `TRUST_BACKFILL`. **Recommended**. Use when you add new trusted columns. This option runs the same processes that the `REVALIDATE` and `RECALCULATE` types run.
     - `REVALIDATE`. Use when you change or add validation rules.
     - `RECALCULATE`. Use when you change trust rules.
     - `TOKENIZE`. Use when you need to run the tokenize process on dirty records, but you are unable to run batch jobs.

- *base object name* is the table name of a base object. Run the command on all base objects tables for which you want to recalculate the best version of the truth (BVT). If you are not sure which tables are affected by the schema update, run the command on all base object tables in the schema.
- *api* specifies which API runs the backfill task. `R` is Read API, `W` is Write API, and `B` is both Read and Write APIs. Use `B`.
- *sequence* is the order to run the backfill task in relationship to other tasks. If you are not sure, use `1` to run the backfill task first.

For example, the following command applies trust backfill for the C_CUSTOMER base object:

```
sip_ant -Dnoprompt=true -noinput add_backfill_task -DprocedureName=TRUST_BACKFILL -
DtableName=C_CUSTOMER -DusageType=B -Dsequence=1
```

8. Run the backfill batch job for each base object.

   You can run the batch backfill job from the Hub Console or by using a service integration framework (SIF) API.

Run From	Steps
Hub Console	1. In the MDM Hub Console, open the Batch Viewer tool. 2. From the Batch Viewer navigation pane, select the base object that you want to backfill. If the backfill batch job does not appear in the batch viewer for the base object, select **Batch Viewer** > **Refresh**. 3. Run the backfill batch job.

API	1. Ensure that the Hub Server is running.
	2. Run one of the following APIs:
	• To backfill all base objects, use the ExecuteBatchBackfillAll API.
	Note: To run the backfill on all records, ensure that the `dirtyOnlyInd` parameter is `false`.
	• To backfill a specified base object, use the ExecuteBatchBackfill API.
	Important: Comment out the rowidObjectTable element in the request.
	
	The following code sample shows an ExecuteBatchBackfill request to backfill records in the C_BO_TRUST base object:
	`<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/" xmlns:urn="urn:siperian.api">`
	`<soapenv:Header/>`
	`<soapenv:Body>`
	`<urn:ExecuteBatchBackfill>`
	`<urn:username><user name></urn:username>`
	`<urn:password>`
	`<urn:password><password></urn:password>`
	`</urn:password>`
	`<urn:orsId>localhost-orclsnl-UTSOURCE</urn:orsId>`
	`<urn:asynchronousOptions>`
	`<urn:isAsynchronous>false</urn:isAsynchronous>`
	`</urn:asynchronousOptions>`
	`<urn:tableName>C_BO_TRUST</urn:tableName>`
	`<!--urn:rowidObjectTable>?</urn:rowidObjectTable-->`
	`<urn:dirtyOnlyInd>false</urn:dirtyOnlyInd>`
	`</urn:ExecuteBatchBackfill>`
	`</soapenv:Body>`
	`</soapenv:Envelope>`

9. **Schema Update with Data Change**. Disable the Oracle GoldenGate mapping for the table C_AGREEMENT and change the mapping to table C_AGREEMENT_XREF_NEW_FROM_A.
    a. Turn off mapping for the entire C_AGREEMENT base object and all related tables by using `sip_ant disable_replication` command:

        ```
        sip_ant -Dnoprompt=true -noinput disable_replication -DtableName=C_AGREEMENT
        ```

    b. Create a mapping from C_AGREEMENT_XREF to C_AGREEMENT_XREF_NEW_FROM_A by using the `sip_ant remap` command. Data in source and target tables must be the same. The `sip_ant remap` command automatically creates the table C_AGREEMENT_XREF_NEW_FROM_A. If the table exists, the command fails. Run the following command:

        ```
        sip_ant -Dnoprompt=true -noinput remap -DtableName=C_AGREEMENT_XREF -
        DmapTableName=C_AGREEMENT_XREF_A
        ```

10. **Schema Update with Data Change**. Start the reload of data to C_AGREEMENT from C_AGREEMENT_XREF. You do not need to wait for the data to reload before you proceed to the next step.
11. If you conduct user acceptance validation, perform the following steps:
    a. Note the system change number (SCN).
    b. Perform user acceptance validation.
    c. Flash back to the noted SCN.
12. To continue the upgrade process, enter the following swing command:

    ```
    sip_ant swing_continue
    ```

    The swing command executes the following processes:
    1. Processes data changes from the passive environment by starting replication replay.
    2. Detects completion of replication catch-up in the passive environment.
    3. Sends a message from the passive environment to the active environment to disallow batch processes in the active environment.
    4. Enables writeable API services in the target system.
    5. Synchronizes sequences. The new sequence values on the target system are higher than on the source system.
    6. Disables writeable SIF services in the active environment.
    7. Complete the replication processing in the passive environment.
13. Verify that the application server is running in the passive environment, and then redirect services from the active environment to the passive environment.

    The passive environment is live for read and write services.
14. **Schema Update with Data Change**. Process the delta on the C_AGREEMENT table.

    The implementation resource writes the delta to handle the data that came from the active environment after you reload the passive environment.

15. **MDM Upgrade** and **Schema Update**. Run the backfill job on dirty records for each base object.

Run From	Steps
Hub Console	1. In the MDM Hub Console, open the Batch Viewer tool. 2. From the Batch Viewer navigation pane, select the base object that you want to backfill. If the backfill batch job does not appear in the batch viewer for the base object, select **Batch Viewer** > **Refresh**. 3. To backfill only on the dirty records, select **For dirty records only**. 4. Run the backfill batch job.

API	1. Ensure that the MDM Hub Server is running.
	2. Run one of the following APIs.
	• To backfill all base objects, use the ExecuteBatchBackfillAll API.
	Note: To run the backfill only on the dirty records, set the `dirtyOnlyInd` parameter to `true`.
	• To backfill a specified base object, use the ExecuteBatchBackfill API.
	
	Important: Comment out the rowidObjectTable element in the request.
	The following code sample shows an ExecuteBatchBackfill request to backfill the dirty records in the C_BO_TRUST base object:
	`<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/" xmlns:urn="urn:siperian.api">`
	`<soapenv:Header/>`
	`<soapenv:Body>`
	`<urn:ExecuteBatchBackfill>`
	`<urn:username><user name></urn:username>`
	`<urn:password>`
	`<urn:password><password></urn:password>`
	`</urn:password>`
	`<urn:orsId>localhost-orclsnl-UTSOURCE</urn:orsId>`
	`<urn:asynchronousOptions>`
	`<urn:isAsynchronous>false</urn:isAsynchronous>`
	`</urn:asynchronousOptions>`
	`<urn:tableName>C_BO_TRUST</urn:tableName>`
	`<!--urn:rowidObjectTable>?</urn:rowidObjectTable-->`
	`<urn:dirtyOnlyInd>true</urn:dirtyOnlyInd>`
	`</urn:ExecuteBatchBackfill>`
	`</soapenv:Body>`
	`</soapenv:Envelope>`

16. Remove the backfill tasks from the C_REPOS_BACKFILL_TASK table. The table must be empty so that other batch jobs can run.

```
Delete from c_repos_zdt_backfill_task;
COMMIT;
```

17. **MDM Upgrade** and **Schema Update**. Run a tokenize batch job on dirty records for each base object.

Run From	Steps
Hub Console	1. In the MDM Hub Console, open the Batch Viewer tool. 2. From the Batch Viewer navigation pane, expand the base object for which you want to regenerate all the match tokens. 3. Expand **Generate Match Tokens**. 4. Select the batch job that you want to use to generate match tokens. 5. Clear **Re-generate All Match Tokens**. 6. Click **Execute Batch**.
API	To run the Generate Match Tokens batch job on dirty records only, use the ExecuteBatchGenerateMatchTokens request with the `fullRestripInd` attribute set to `0`. The following code sample shows an ExecuteBatchGenerateMatchTokens request to create match tokens for dirty records in the C_PARTY base object: `<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/" xmlns:urn="urn:siperian.api"> <soapenv:Header/> <soapenv:Body> <urn:executeBatchGenerateMatchTokens> <urn:username>admin</urn:username> <urn:password> <urn:password>admin</urn:password> <urn:encrypted>false</urn:encrypted> </urn:password> <urn:orsId>localhost-orclsnl-UTSOURCE</urn:orsId> <urn:asynchronousOptions> <urn:isAsynchronous>false</urn:isAsynchronous> <urn:jmsReplyTo></urn:jmsReplyTo> <urn:jmsCorrelationId></urn:jmsCorrelationId> </urn:asynchronousOptions> <urn:tableName>C_PARTY</urn:tableName> <urn:fullRestripInd>0</urn:fullRestripInd> </urn:executeBatchGenerateMatchTokens> </soapenv:Body> </soapenv:Envelope>`

18. Run the following command to finalize the upgrade process:

```
sip_ant swing_finalize
```

The swing command executes the following processes:

1. Enables batch services in the target system.
2. Removes the ZDT event queue from the following tables:
   - C_REPOS_ZDT_EVENT_QUEUE
   - C_REPOS_ZDT_REPLICAT_EXCEPTION
3. Updates the following tables and sets their values:
   a. C_REPOS_ZDT_ENV_STATE; set state = NULL, state_ts = NULL, state_desc = NULL, updated_by=NULL, update_date=NULL
   b. C_REPOS_ZDT_STATUS; set REPLICATION_TARGET_IND = 0
   c. C_REPOS_ZDT_STATUS; set ACTIVE_UPGRADE_IND = 0
4. Configures sequences on the active environment to be even.
5. Undeploys ZDT.

6. Configures and starts the extract on the active environment, which is the new source.

19. **Schema Update with Data Change**. If any user data foreign keys change during the update, run foreign key validation. Run the ExecuteBatchValidateFKRelationships SIF API for each base object.

    Any violations that came from the active environment after the lookup data is updated are detected.

20. **Schema Update**. If violations are detected, fix the violations.

    Tip: If the violations are not severe, you can fix the violations after you complete the upgrade. If you cannot fix some violations, contact Informatica Global Customer Support.

21. Get the current SCN from the passive environment:

```
SQL-CMX_ORS_B> select current_scn from v$database;

CURRENT_SCN
-----------
    2880593
```

22. Export the passive environment using data pump with the SCN:

```
c:> <ors username>/<password>@<tns entry name>
    directory=<DATA_PUMP_DIR_OBJECT>
    dumpfile=<mrm_backup_envb.dmp>
    logfile=<mrm_backup_after_upgrade.log>
    parallel=8
    job_name=<EXPORT_AFTER_UPGRADE>
    flashback_scn=<CURRENT_SCN from the previous step>
```

23. Resume batch jobs in the passive environment.

Upgrade Steps Controlled from the Active Environment

The active environment must be prepared after you complete the upgrade steps on the passive environment. Run the steps from the active environment.

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.

Before you begin, open the following repository tables and make a note of the values in the columns.

C_REPOS_ZDT_STATUS

Record the values for all columns. You need these values in step 6.

C_REPOS_DB_RELEASE

Record the values for the following columns. You need these values in step 7.

- db_password, tns_name
- connection_port
- oracle_sid
- database_host

- connect_url
- database_id
- connection_type
- proxy_ind
- db_proxy_username
- db_proxy_password
- db_replication_username
- db_replication_password
- debug_ind
- debug_level
- debug_file_name
- debug_file_pat

1. **Infrastructure Upgrade**. Upgrade hardware and third-party software in the active environment.
2. Stop the application servers, close connections, and then drop and re-create the schema in the active environment.
   a. Stop the application servers in the active environment.
   b. Close all the connections, such as SQL*Plus, Toad, and application server, to the active environment.
   c. Drop the schema in active environment, and then create it.

```
Drop schema A - using system user (sqlplus system/password@tnsname)
 SQL> drop user envA cascade;

Create schema A - using system user (sqlplus system/password@tnsname)
 SQL> <hub_server_install>/resources/database/custom_scripts/oracle/import/
@mk_cmx_ors_user; -- supply the schema name as A
```

3. Re-create the Operational Reference Store by importing the dump file that was generated from the passive environment.

```
C:\> impdp <dba_username>/<dba_password>@<tns_entry_name>
     directory=<DATA_PUMP_DIR_OBJECT>
     dumpfile==<mrm_backup_envb.dmp>
     logfile=<mrm_restore_after_upgrade.log>
     content=all
     remap_schema=<from_user>:<to_user>
     parallel=8
     job_name=<RESTORE_ENVB>
```

   When the schema is being created, you can safely ignore the following messages:

```
ORA-39083: Object type TYPE failed to create with error:
ORA-02304: invalid object identifier literal
The type already exists, and therefore it is not re-created
```

4. Remove event queues and drop tables.

```
/* 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;
```

5. Run the following command to install the Java utility file:

```
sip_ant install_utility
```

6. To complete the swing installation process, run the following command:

```
sip_ant swing_finish
```

The replication parameter files install.

7. Check the environment-specific settings on the active environment in the C_REPOS_DB_RELEASE table.

   All the values in the table must be local and the database entries must point to the local database. If necessary, update the values to match the values that were in the C_REPOS_DB_RELEASE table before you dropped the schema.

8. **MDM Upgrade**. On the new passive environment, upgrade the Multidomain MDM and configure the Process Server.

CHAPTER 3

Troubleshooting

This chapter includes the following topics:

Batch job fails when backfill tasks are registered

If you run batch jobs with backfill tasks enabled, an error occurs and the batch job fails.

The MDM Hub does not allow batch jobs to run in a target replication environment while backfill tasks are registered. The MDM Hub does not allow batch jobs to run because of the high likelihood of locking conflicts when backfill batch jobs are running concurrently with regular batch jobs. After the backfill task has completed and the backfill tasks are de-registered, you can run batch jobs.

Reset Matches on Target

The Reset Match logic is not replicated to the target environment. MET migration will not reset the match table. You must do this using a script. Tokenize backfill does not impact the match table. After a match-related change list is applied, a change must be made to the target match rules in order to trigger Reset Match.

Replication is not working

If the deploy_zdt call does not complete, the ZDT replication might not work between the source and target databases.

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 E_ENVA
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. 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' );
   ```

3. 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

INDEX