Oracle® Fusion Middleware Using Oracle GoldenGate for Heterogeneous Databases





Oracle Fusion Middleware Using Oracle GoldenGate for Heterogeneous Databases, 19c (19.1.0)

E98067-21

Copyright © 2017, 2024, Oracle and/or its affiliates.

Primary Author: Oracle Corp.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

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

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

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

	Prerace	
	Audience Documentation Accessibility	xiv Xiv
	Related Information	XiV
	Conventions	ΧİV
Part	What is Oracle GoldenGate for Heterogeneous Databases?	
Part	Using Oracle GoldenGate for DB2 LUW	
1	Understanding What's Supported for DB2 LUW	
	Supported DB2 LUW Data Types	1-1
	Non-Supported DB2 LUW Data Types	1-2
	Supported Objects and Operations for DB2 LUW	1-2
	Non-Supported Objects and Operations for DB2 LUW	1-4
	System Schemas	1-4
	Supported Object Names	1-4
2	Preparing the System for Oracle GoldenGate	
	Configuring the Transaction Logs for Oracle GoldenGate	2-1
	Retaining the Transaction Logs	2-1
	Specifying the Archive Path	2-2
	Preparing Tables for Processing	2-2
	Triggers and Cascade Constraints Considerations	2-3
	Assigning Row Identifiers	2-3
	How Oracle GoldenGate Determines the Kind of Row Identifier to Use	2-3
	Using KEYCOLS to Specify a Custom Key	2-4
	Preventing Key Changes	2-4
	Enabling Change Capture	2-4
	Maintaining Materialized Query Tables	2-5



	Database Configuration for DB2 LUW	2-5
	Database User for Oracle GoldenGate Processes for DB2 LUW	2-6
	Setting the Session Character Set	2-6
	Preparing for Initial Extraction	2-7
	Specifying the DB2 LUW Database in Parameter Files	2-7
3	Configuring Oracle GoldenGate for DB2 LUW	
	What to Expect from these Instructions	3-1
	Where to Get More Information	3-1
	Configuring the Primary Extract	3-2
	Configuring the Data Pump Extract	3-3
	Creating a Temporal Table	3-4
	Support for Temporal Tables	3-4
	Replicating with Temporal Tables	3-4
	Converting	3-5
	Creating a Checkpoint Table	3-9
	Configuring the Replicat Parameter File	3-9
	When to Start Replicating Transactional Changes	3-10
	Testing Your Configuration	3-10
4	Understanding What's Supported for DB2 for i	
	Supported DB2 for i Data Types	4-1
	Non-Supported DB2 for i Data Types	4-2
	Supported Objects and Operations for DB2 for i	4-2
	Non-Supported Objects and Operations for DB2 for i	4-3
	Oracle GoldenGate Parameters Not Supported for DB2 for i	4-3
	Supported Object Naming Conventions	4-3
	System Schemas for DB2 for i	4-4
	Supported Character Sets	4-4
5	Preparing the System for Oracle GoldenGate	
	Preparing the Journals for Data Capture by Extract	5-1
	Allocating Journals to an Extract Group	5-1
	Setting Journal Parameters	5-1
	Deleting Old Journal Receivers	5-2
	Specifying Object Names	5-3
	Preparing Tables for Processing	5-3



	Assigning Row Identifiers	5-3
	How Oracle GoldenGate Determines the Kind of Row Identifier to Use	5-4
	Using KEYCOLS to Specify a Custom Key	5-4
	Preventing Key Changes	5-4
	Disabling Constraints on the Target	5-5
	Enabling Change Capture	5-5
	Specifying a Default Journal	5-6
	Removing a Default Journal Specification	5-6
	Maintaining Materialized Query Tables	5-6
	Specifying the Oracle GoldenGate Library	5-6
	Adjusting the System Clock	5-6
	Configuring the ODBC Driver	5-7
	Configuring ODBC on Linux	5-7
	Configuring ODBC on Windows	5-8
6	Configuring Oracle GoldenGate for DB2 for i	
	What to Expect from this Procedure	6-1
	Getting Started with Oracle GoldenGate	6-1
	Creating the Oracle GoldenGate Instance	6-2
	Creating a GLOBALS File	6-2
	Creating a Data Definitions File	6-2
	Encrypting the Extract and Replicat Passwords	6-2
	Configuring Extract for Change Capture from DB2 for i	6-3
	Configuring the Primary Extract	6-3
	Configuring the Data Pump	6-4
	Configuring Replicat for Change Delivery to DB2 for i	6-5
	Creating a Checkpoint Table	6-5
	Configuring Replicat	6-6
	Next Steps in the Deployment	6-7
	When to Start Replicating Transactional Changes	6-7
	Starting Extract During Instantiation	6-8
	Changing the Position of Extract to a Later Time	6-8
	Testing Your Configuration	6-8
7	Instantiating and Starting Oracle GoldenGate Replication	
	About the Instantiation Process	7-1
	Overview of Basic Oracle GoldenGate Instantiation Steps	7-2
	Satisfying Prerequisites for Instantiation	7-2
	Configure Change Capture and Delivery	7-2
	Add Collision Handling	7-2



	Prepare the Target Tables	7-3
	Making the Instantiation Procedure More Efficient	7-3
	Share Parameters Between Process Groups	7-3
	Use Parallel Processes	7-3
	Configuring the Initial Load	7-3
	Configuring an Initial Load from File to Replicat	7-4
	Configuring an initial load with a database utility	7-7
	Adding Change-Capture and Change-Delivery processes	7-7
	Add the Primary Extract	7-8
	Understanding the Primary Extract Start Point	7-8
	Establishing the Required and Optional Extract Start Points	7-8
	Add the Local Trail	7-10
	Add the Data Pump Extract Group	7-10
	Add the Remote Trail	7-10
	Add the Replicat Group	7-11
	Performing the Target Instantiation	7-11
	To Perform Instantiation from File to Replicat	7-11
	To Perform Instantiation with a Database Utility	7-12
	Monitoring Processing after the Instantiation	7-14
	Backing up Your Oracle GoldenGate Environment	7-14
	Positioning Extract After Startup	7-14
8	Using Remote Journal	
	Preparing to Use Remote Journals	8-1
	Adding a Remote Journal	8-2
	What Happens During Add Remote Journal Processing?	8-2
	Guidelines For Adding a Remote Journal	8-3
Part	IV Using Oracle GoldenGate for DB2 for z/OS	
9	Understanding What's Supported for DB2 for z/OS	
	Supported DB2 for z/OS Data Types	9-1
	Non-Supported DB2 for z/OS Data Types	9-1
	Supported Objects and Operations for DB2 z/OS	9-2
	Non-Supported Objects and Operations for DB2 for z/OS	9-2
10	Preparing the DB2 for z/OS Database for Oracle GoldenGate	
	Preparing Tables for Processing	10-1
	Disabling Triggers and Cascade Constraints	10-1



	Assigning Row Identifiers	10-1
	How Oracle GoldenGate Determines the Kind of Row Identifier to Use	10-2
	Using KEYCOLS to Specify a Custom Key	10-2
	Handling ROWID Columns	10-2
	Configuring a Database Connection	10-3
	Database Configuration for DB2 z/OS	10-3
	Database User for Oracle GoldenGate Processes	10-3
	Setting Initialization Parameters	10-4
	Specifying the Path to the Initialization File	10-5
	Ensuring ODBC Connection Compatibility	10-5
	Specifying the Number of Connection Threads	10-5
	Monitoring Processes	10-6
	Interpreting Statistics for Update Operations	10-6
	Supporting Globalization Functions	10-6
	Replicating From a Source that Contains Both ASCII and EBCDIC	10-6
	Specifying Multi-Byte Characters in Object Names	10-7
11	Preparing the DB2 for z/OS Transaction Logs for Oracle GoldenGate	
	Making Transaction Data Available	11-1
	Enabling Change Capture	11-1
	Enabling Access to Log Records	11-1
	Sizing and Retaining the Logs	11-2
	Using Archive Logs on Tape	11-3
	Controlling Log Flushes	11-3
Part	V Using Oracle GoldenGate for MySQL	
12	Understanding What's Supported for MySQL	
	Character Sets in MySQL	12-1
	Supported MySQL Data Types	12-2
	Limitations and Clarifications	12-2
	Non-Supported MySQL Data Types	12-3
	Supported Objects and Operations for MySQL	12-4
	Non-Supported Objects and Operations	12-5
	System Schemas	12-5
13	Preparing and Configuring the System for Oracle GoldenGate	
	Database User for Oracle GoldenGate Processes for MySQL	13-1
	Ensuring Data Availability	13-3



	Setting Logging Parameters	13-3
	Adding Host Names	13-5
	Setting the Session Character Set	13-5
	Preparing Tables for Processing	13-5
	Assigning Row Identifiers	13-6
	How Oracle GoldenGate Determines the Kind of Row Identifier to Use	13-6
	Tables with a Primary Key Derived from a Unique Index	13-6
	How to Specify Your Own Key for Oracle GoldenGate to Use	13-7
	Limiting Row Changes in Tables That Do Not Have a Key	13-7
	Triggers and Cascade Constraints Considerations	13-7
	Changing the Log-Bin Location	13-8
	Configuring Bi-Directional Replication	13-8
	Configuring MySQL for Remote Capture	13-9
	Configuring a Two-way SSL Connection in MySQL Capture and Delivery	13-10
	Capturing using a MySQL Replication Slave	13-11
	Other Oracle GoldenGate Parameters for MySQL	13-11
	Positioning Extract to a Specific Start Point	13-14
L4	Using DDL Replication	
	Plug-in Based DDL Configuration Prerequisites and Considerations	14-1
	Installing DDL Replication	14-2
	Using the Metadata Server	14-3
	Using DDL Filtering for Replication	14-3
	Troubleshooting Plug-in Based DDL Replication	14-5
	Uninstalling Plug-In Based DDL Replication	14-5
Part	VI Using Oracle GoldenGate for PostgreSQL	
L5	Understanding What's Supported for PostgreSQL	
	Supported Databases	15-1
	Details of Supported PostgreSQL Data Types	15-1
	Supported PostgreSQL Data Types	15-1
	Non-Supported PostgreSQL Data Types	15-3
	Supported Objects and Operations for ProtgreSQL	15-4
	Tables and Views	15-5
	Sequences and Identity Columns	15-5



Preparing the Database for Oracle GoldenGate

	Database Configuration	16-1
	Database Settings for Azure Database for PostgreSQL, Amazon Aurora, and Amazon RDS	16-2
	Establishing Oracle GoldenGate Credentials	16-2
	Assigning Credentials to Oracle GoldenGate	16-3
	Securing the Oracle GoldenGate Credentials	16-5
	Configuring a Database Connection	16-5
	Configuring a Database Connection in Linux	16-6
	Configuring SSL Support for PostgreSQL	16-7
	Preparing Tables for Processing	16-7
	Disabling Triggers and Cascade Constraints on the Target	16-7
	Ensuring Row Uniqueness for Tables	16-8
	Enabling Table-Level Supplemental Logging	16-8
L7	Configuring Extract	
	About Extract	17-1
	Initial Load Extract	17-1
	Change Data Capture Extract	17-1
	Extract Deployment Options	17-1
	Prerequisites for Creating an Extract	17-2
	Registering an Extract for PostgreSQL	17-2
	Creating a Change Data Capture Extract	17-2
	Modifying DDL for Extract	17-4
L8	Configuring Replicat	
	About Replicat	18-1
	Replicat Deployment Options	18-1
	Prerequisites for Creating a Replicat	18-1
	Creating a Checkpoint Table	18-2
	Creating a Replicat	18-2
L9	Additional Considerations	
	Adding a Heartbeat Table	19-1
	Running the Heartbeat Update and Purge Function	19-1
	Enabling Bi-Directional Loop Detection	19-2
	Deleting an Extract	19-3
	Removing Table-level Supplemental Logging	19-3



Part VII Using Oracle GoldenGate for SQL Server

20	Understanding What's Supported for SQL Server	
	Supported Objects and Operations for SQL Server	20-1
	Non-Supported Objects and Operations for SQL Server	20-1
	Requirements for Table Level DDL Changes	20-2
	Supported SQL Server Data Types	20-3
	System Schemas for SQL Server	20-5
	Non-Supported SQL Server Data Types and Features	20-5
21	Preparing the System for Oracle GoldenGate	
	Prepare Database Users and Privileges	21-1
	Oracle GoldenGate for SQL Server	21-1
	Amazon RDS User Permissions and Requirements	21-2
	Azure SQL Database	21-3
	Configuring a Database Connection	21-4
	Configuring an Extract Database Connection	21-4
	Configuring a Replicat Database Connection	21-4
	Using ODBC or Default OLE DB	21-4
	Using OLE DB with USEREPLICATIONUSER	21-5
	Configuring a Database Connection on Linux	21-6
	Configuring a Database Connection on Windows	21-7
	Preparing Tables for Processing	21-8
	Disabling Triggers and Cascade Constraints on the Target	21-8
	Assigning Row Identifiers	21-9
	How Oracle GoldenGate Determines the Kind of Row Identifier to Use	21-9
	Using KEYCOLS to Specify a Custom Key	21-9
	Improving IDENTITY Replication with Array Processing	21-9
	Globalization Support	21-10
22	Preparing the Database for Oracle GoldenGate — CDC Capture	
	Enabling CDC Supplemental Logging	22-1
	Purging CDC Staging Data	22-3
	Enabling Bi-Directional Loop Detection	22-4
23	Requirements Summary for Capture and Delivery of Databases in ar Always On Availability Group	l
	Database Connection	23-1



	Supplemental Logging	23-2
C	Operational Requirements and Considerations	23-2
24	CDC Capture Method Operational Considerations	
T	uning SQL Server Change Data Capture	24-1
C	Dracle GoldenGate CDC Object Versioning	24-2
V	/alid and Invalid Extract Parameters for SQL Server Change Data Capture	24-2
D	Details of the Oracle GoldenGate CDC Cleanup Process	24-3
C	Changing from Classic Extract to a CDC Extract	24-5
R	Restoring a Source Database Keeping CDC Data	24-5
Part V	/III Using Oracle GoldenGate with Sybase	
25 L	Understanding What's Supported for Sybase	
S	Supported Sybase Data Types	25-1
	Integers	25-1
	Floating-Point Numbers	25-2
	Character Data	25-2
	Dates and Timestamps	25-3
	Large Objects	25-3
	Money Types	25-4
	IDENTITY Type	25-4
	text, image, and unitext Data Types	25-4
	User-Defined Data Types	25-4
N	Non-Supported Sybase Data Types	25-4
S	Supported Operations and Objects for Sybase	25-5
Ν	Non-Supported Operations and Objects for Sybase	25-6
26 F	Preparing the System for Oracle GoldenGate	
	Database Configuration	26-1
	Database User for Oracle GoldenGate Processes	26-1
P	Preparing Tables for Processing	26-2
	Disabling Triggers and Cascade Constraints	26-2
	Assigning Row Identifiers	26-3
	Limiting Row Changes in Tables that do not Have a Key	26-3
	Replicating Encrypted Data	26-3
Р	Preparing the Transaction Logs	26-3
	Initializing the Secondary Truncation Point	26-4
	Sizing and Retaining the Logs	26-4



Part IX Using Oracle GoldenGate for Teradata

Understanding What's Supported for Teradata	
Supported Teradata Data Types	27-1
Limitations of Support for Numeric Data Types	27-2
Limitations of Support for Single-byte Character Data Types	27-3
Conditions and Limitations of Support for Multi-byte Character Data	27-3
Limitations of Support for Binary Data Types	27-3
Limitations of Support for Large Object Data Types	27-3
Limitations of Support for Date Data Types	27-3
Limitations of Support for IDENTITY Data Types	27-4
Supported Objects and Operations for Teradata	27-4
Non-Supported Operations for Teradata	27-4
Preparing the System for Oracle GoldenGate	
Preparing Tables for Processing	28-1
Disabling Triggers and Cascade Constraints	28-1
Assigning Row Identifiers	28-2
How Oracle GoldenGate Determines the Kind of Row Identifier to Use	28-2
Using KEYCOLS to Specify a Custom Key	28-2
ODBC Configuration for Teradata	28-2
Database User for Oracle GoldenGate Processes for Teradata	28-2
Configuring Oracle GoldenGate	
Creating a Checkpoint Table	29-1
Configuring Oracle GoldenGate Replicat	29-1
Additional Oracle GoldenGate Configuration Guidelines	29-2
Handling Massive Update and Delete Operations	29-2
Preventing Multiple Connections	29-2
Performing Initial Synchronization	29-2
Common Maintenance Tasks	
Modifying Columns of a Table	30-1



Part X Using Oracle GoldenGate for Oracle TimesTen

Understanding What's Supported for Oracle TimesTen	
Supported Objects and Operations for TimesTen	31-1
Non-Supported TimesTen Data Types and Features	31-1
Supported TimesTen Data Types	31-1
Limitations and Non-supported Items for Oracle TimesTen	31-2
Preparing the System for Oracle GoldenGate	
Export All Variables	32-1
Configuring the TimesTen Client ODBC Driver	32-1
Configuring ODBC on Linux	32-1
Preparing Tables for Processing	32-2
Disabling Triggers and Cascade Constraints	32-3
Assigning Row Identifiers	32-3
Configuring Oracle GoldenGate	32-3
Configuring Oracle GoldenGate Replicat	32-3
Additional Oracle GoldenGate Configuration Guidelines	32-4



Preface

This guide helps you get started with using Oracle GoldenGate on heterogeneous database systems supported with this release.

Topics:

- Audience
- Documentation Accessibility
- Related Information
- Conventions

Audience

Using Oracle GoldenGate for Heterogeneous Databases is intended for database and system administrators who are responsible for implementing Oracle GoldenGate and managing the databases for an organization.

Documentation Accessibility

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

Access to Oracle Support

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

Related Information

The Oracle GoldenGate Product Documentation Libraries are found at

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

Additional Oracle GoldenGate information, including best practices, articles, and solutions, is found at:

Oracle GoldenGate A-Team Chronicles

Conventions

The following text conventions are used in this document:



Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, such as "From the File menu, select Save ." Boldface also is used for terms defined in text or in the glossary.
italic italic	Italic type indicates placeholder variables for which you supply particular values, such as in the parameter statement: TABLE table_name. Italic type also is used for book titles and emphasis.
monospace MONOSPACE	Monospace type indicates code components such as user exits and scripts; the names of files and database objects; URL paths; and input and output text that appears on the screen. Uppercase monospace type is generally used to represent the names of Oracle GoldenGate parameters, commands, and user-configurable functions, as well as SQL commands and keywords.
UPPERCASE	Uppercase in the regular text font indicates the name of a utility unless the name is intended to be a specific case.
{}	Braces within syntax enclose a set of options that are separated by pipe symbols, one of which must be selected, for example: {option1 option2 option3}.
[]	Brackets within syntax indicate an optional element. For example in this syntax, the SAVE clause is optional: CLEANUP REPLICAT group_name [, SAVE count]. Multiple options within an optional element are separated by a pipe symbol, for example: [option1 option2].



Part I

What is Oracle GoldenGate for Heterogeneous Databases?

Oracle GoldenGate is a comprehensive software package for real-time data capture and replication in heterogeneous IT environments.

The product set enables high availability solutions, real-time data integration, transactional change data capture, data replication, transformations, and verification between operational and analytical enterprise systems. Oracle GoldenGate brings extreme performance with simplified configuration and management, support for cloud environments, expanded heterogeneity, and enhanced security.

You can use the following supported heterogeneous databases with Oracle GoldenGate.

- DB2 LUW
- DB2 for i
- DB2 z/OS
- MySQL
- PostgreSQL
- SQL Server
- Teradata
- TimesTen

Each database that Oracle GoldenGate supports has it's own requirements and configuration.

Many databases contain schemas that are intended for use by the database or the system. In most cases, objects in these schemas should not be automatically replicated through the use of a wildcard specification. For the databases and schemas listed in this book, there are specific schemas or objects that will not be automatically replicated by Oracle GoldenGate unless they are explicitly specified without a wildcard.

This book is divided into parts so that you can easily find information that is relevant to your environment. See *Installing Oracle GoldenGate* for system requirements and installation details for each of these databases.



Part II

Using Oracle GoldenGate for DB2 LUW

With Oracle GoldenGate for DB2 LUW, you can perform initial loads and capture transactional data from supported DB2 LUW database versions and replicate the data to a DB2 LUW database or other supported Oracle GoldenGate targets, such as an Oracle Database.

Oracle GoldenGate for DB2 LUW supports data filtering, mapping, and transformations unless noted otherwise in this documentation.

This part describes tasks for configuring and running Oracle GoldenGate for DB2 LUW.

- Understanding What's Supported for DB2 LUW
 This chapter contains information on database and table features supported by Oracle GoldenGate for DB2 LUW.
- Preparing the System for Oracle GoldenGate
- Configuring Oracle GoldenGate for DB2 LUW



1

Understanding What's Supported for DB2 LUW

This chapter contains information on database and table features supported by Oracle GoldenGate for DB2 LUW.

Topics:

- Supported DB2 LUW Data Types
- Non-Supported DB2 LUW Data Types
- Supported Objects and Operations for DB2 LUW
- Non-Supported Objects and Operations for DB2 LUW
- System Schemas
- Supported Object Names

Supported DB2 LUW Data Types

Oracle GoldenGate supports all DB2 LUW data types, except those listed in Non-Supported DB2 LUW Data Types.

Limitations of Support

Oracle GoldenGate has the following limitations for supporting DB2 LUW data types:

- Oracle GoldenGate supports multi-byte character data types and multi-byte data stored in character columns. Multi-byte data is only supported in a like-to-like configuration.
 Transformation, filtering, and other types of manipulation are not supported for multi-byte character data.
- BLOB and CLOB columns must have a logged clause in their definitions.
- Due to limitations in the IBM DB2READLOG interface, Oracle GoldenGate does not support coordination of transactions across nodes in a DB2 Database Partitioning Feature (DPF) environment. In DPF, a transaction may span multiple nodes, depending upon how the data is partitioned.

However, if you need to capture from it, you can do it with certain limitations. Check the Oracle Support note Does Oracle GoldenGate Support DB2 LUW Data Partitioning Feature (DPF)? (DocID 2763006.1)

- GRAPHIC and VARGRAPHIC columns must be in a database, where the character set is UTF16. Any other character set causes the Oracle GoldenGate to abend.
- The support of range and precision for floating-point numbers depends on the host machine. In general, the precision is accurate to 16 significant digits, but you should review the database documentation to determine the expected approximations. Oracle GoldenGate rounds or truncates values that exceed the supported precision.

- Extract fully supports the capture and apply of TIMESTAMP (0) through TIMESTAMP (12) when the output trail format is 19.1 or higher. Otherwise Extract truncates data from TIMESTAMP (10) through TIMESTAMP (12) to nanoseconds (maximum of nine digits of fractional time) and issues a warning to the report file.
- Oracle GoldenGate supports timestamp data from 0001/01/03:00:00:00 to 9999/12/31:23:59:59. If a timestamp is converted from GMT to local time, these limits also apply to the resulting timestamp. Depending on the timezone, conversion may add or subtract hours, which can cause the timestamp to exceed the lower or upper supported limit.
- Oracle GoldenGate does not support the filtering, column mapping, or manipulation of large objects that are larger than 4K. You can use the full Oracle GoldenGate functionality for objects that are 4K or smaller.
- Replication of XML columns between source and target databases with the same character set is supported. If the source and target database character sets are different, then XML replication may fail with a database error because some characters may not be recognized (or valid) in the target database character set.

Non-Supported DB2 LUW Data Types

The non-supported DB2 LUW data types are:

- XMLType
- User-defined types
- Negative dates

Supported Objects and Operations for DB2 LUW

Object and operations that are supported for DB2 LUW are:

- Oracle GoldenGate Extract supports cross-endian capture where the database and Oracle GoldenGate are running on different byte order servers. The byte order is detected automatically for DB2 LUW version 10.5 or higher. If the DB2 database auto-detection on the DB2 LUW 10.5 database is not required then you can override it by specifying the TRANLOGOPTIONS MIXEDENDIAN [ON|OFF] parameter. For DB2 LUW version 10.1, this parameter must be used in the Extract parameter file for cross-endian capture. See TRANLOGOPTIONS in Reference for Oracle GoldenGate.
- DB2 pureScale environment is supported.
- Oracle GoldenGate supports the maximum number of columns and column size per table that is supported by the database.
- TRUNCATE TABLE.
- Multi-Dimensional Clustered Tables (MDC).
- Materialized Query Tables. Oracle GoldenGate does not replicate the MQT itself, but only
 the base tables. The target database automatically maintains the content of the MQT
 based on the changes that are applied to the base tables by Replicat.
- Tables with ROW COMPRESSION. In DB2 LUW version 10.1 and later, COMPRESS YES STATIC
 is supported and COMPRESS YES ADAPTIVE are supported.
- Extended row size feature is enabled by default. It is supported with a workaround using
 FETCHCOLS. For any column values that are VARCHAR or VARGRAPHIC data types and are
 stored out of row in the database, you must fetch these extended rows by specifying these



columns using the FETCHCOLS option in the TABLE parameter in the extract parameter file. With this option set, when the column values are out of row then Oracle GoldenGate will fetch its value. If the value is out of and FETCHCOLS is *not* specified then Extract will abend to prevent any data loss. If you do not want to use this feature, set the <code>extended_row_size</code> parameter to <code>DISABLE</code>.

Extended row size feature is enabled, by default. It is supported with a workaround using FETCHCOLS for DB2 LUW 10.1. For any column values that are VARCHAR or VARGRAPHIC data types and are stored out of row in the database, you must fetch these extended rows by specifying these columns using the FETCHCOLS option in the TABLE parameter in the Extract parameter file. With this option set, when the column values are out of row, then Oracle GoldenGate fetches its value. If the value is out of and FETCHCOLS is not specified then Extract abends to prevent any data loss. If you do not want to use this feature, set the extended_row_size parameter to DISABLE. For DB2 LUW 10.5 and higher out of row values are captured seamlessly by Extract. FETCHCOLS is no more needed to capture out of row columns from these database versions.

- Temporal tables with DB2 LUW 10.1 FixPack 2 and greater are supported. This is the default for Replicat.
- Supported options with SHOWTRANS

```
SHOWTRANS [transaction_ID] [COUNT n]
[DURATION duration unit]
[TABULAR][FILE file name] |
```

Options with SKIPTRANS and FORCETRANS.

```
SKIPTRANS transaction_ID
[FORCE] FORCETRANS transaction ID [FORCE]
```

- Limitations on Automatic Heartbeat Table support are as follows:
 - [THREAD n] [DETAIL] is not supported.
 - Oracle GoldenGate heartbeat parameters frequency and purge frequency are accepted in seconds and days. However, the DB2 LUW task scheduler accepts its schedule only in cron format so the Oracle GoldenGate input value to cron format may result in some loss of accuracy. For example:

```
ADD HEARTBEATTABLE, FREQUENCY 150, PURGE FREQUENCY 20
```

This example sets the FREQUENCY to 150 seconds, which is converted to the closest minute value of 2 minutes, so the heartbeat table is updated every 120 seconds instead of every 150 seconds. Setting PURGE_FREQUENCY to 20 means that the history table is purged at midnight on every 20th day.

- The following are steps are necessary for the heartbeat scheduled tasks to run:
 - 1. Set the DB2_ATS_ENABLE registry variable to db2set DB2_ATS_ENABLE=YES.
 - Create the SYSTOOLSPACE tablespace if it does not already exist:

```
CREATE TABLESPACE SYSTOOLSPACE IN IBMCATGROUP MANAGED BY AUTOMATIC STORAGE EXTENTSIZE 4
```

3. Ensure instance owner has Database administration authority (DBADM):

```
GRANT DBADM ON DATABASE TO instance owner name
```



Non-Supported Objects and Operations for DB2 LUW

Objects and operations for DB2 LUW that are not supported by Oracle GoldenGate are:

- Schema, table or column names that have trailing spaces
- Multiple instances of a database
- Datalinks
- Extraction or replication of DDL (data definition language) operations
- Generated columns (GENERATE ALWAYS clause)



• DB2 Data Partitioning Feature (DPF) is not supported. DPF doesn't provide a way to read the log records in a coordinated fashion across all the nodes in a partition. So, there is no way to ensure that even if all nodes are being replicated with separate Extracts, it would be possible to ensure that all records from all transactions are applied in the correct temporal order. This may occur due to a number of factors including caching in the database and the underlying operating system not allowing the Extracts to have visibility into whether there are any cached and not yet visible entries that may affect the ordering of the record operations across all partitions. Due to this uncertainty, it is not possible to ensure that transactions that span partitions, or primary key updates can be replicated in a consistent manner.

System Schemas

The following schemas or objects are not be automatically replicated by Oracle GoldenGate unless they are explicitly specified without a wildcard.

- "SYSIBM"
- "SYSCAT"
- "SYSSTAT"
- "SYSPROC"
- "SYSFUN"
- "SYSIBMADMIN"
- "SYSTOOLS"
- "SYSPUBLIC"

Supported Object Names

For a list of characters that are supported in object names, see Supported Database Object Names in *Administering Oracle GoldenGate*.



Preparing the System for Oracle GoldenGate

This chapter describes how to prepare the environment to run Oracle GoldenGate on DB2 LUW.

Topics:

- Configuring the Transaction Logs for Oracle GoldenGate
- Preparing Tables for Processing
- Database Configuration for DB2 LUW
- Database User for Oracle GoldenGate Processes for DB2 LUW
- Setting the Session Character Set
- Preparing for Initial Extraction
- Specifying the DB2 LUW Database in Parameter Files

Configuring the Transaction Logs for Oracle GoldenGate

To capture DML operations, Oracle GoldenGatereads the DB2 LUW online logs by default. However, it reads the archived logs if an online log is not available. To ensure the continuity and integrity of Oracle GoldenGateprocessing, configure the logs as follows.

- · Retaining the Transaction Logs
- Specifying the Archive Path

Retaining the Transaction Logs

Configure the database to retain the transaction logs for roll forward recovery by enabling one of the following parameter sets, depending on the database version.

DB2 LUW 9.5 and later:

Set the LOGARCHMETH parameters as follows:

- Set LOGARCHMETH1 to LOGRETAIN.
- Set LOGARCHMETH2 to OFF.

Alternatively, you can use any other LOGARCHMETH options, as long as forward recovery is enabled. For example, the following is valid:

- Set LOGARCHMETH1 to DISK.
- Set LOGARCHMETH2 to TSM.

To determine the log retention parameters:

Connect to the database.

db2 connect to database user username using password

Get the database name.

```
db2 list db directory
```

3. Get the database configuration for the database.

```
db2 get db cfg for database
```

The fields to view are:

```
Log retain for recovery status = RECOVERY User exit for logging status = YES
```

To set the log retention parameters:

1. Issue one of the following commands.

```
To enable USEREXIT:
```

```
db2 update db cfg for database using USEREXIT ON
```

If not using USEREXIT, use this command:

```
db2 update db cfg for database using LOGRETAIN RECOVERY
```

To set LOGARCHMETH:

```
db2 update db cfg for database using LOGARCHMETH1 LOGRETAIN db2 update db cfg for database using LOGARCHMETH2 OFF
```

2. Make a full backup of the database by issuing the following command.

```
db2 backup db database to device
```

3. Place the backup in a directory to which DB2 LUW has access rights. If you get the following message, contact your systems administrator:

```
SQL2061N An attempt to access media "device" is denied.
```

Specifying the Archive Path

Set the DB2 LUW OVERFLOWLOGPATH parameter to the archive log directory. The node attaches automatically to the path variable that you specify.

```
db2 connect to database
db2 update db cfg using overflowlogpath "path"
```

Exclude the node itself from the path. For example, if the full path to the archive log directory is /sdb2logarch/oltpods1/archive/OLTPODS1/NODE0000, then the OVERFLOWLOGPATH value should be specified as /sdb2logarch/oltpods1/archive/OLTPODS1.

Preparing Tables for Processing

The following table attributes must be addressed in an Oracle GoldenGate environment.

- Triggers and Cascade Constraints Considerations
- Assigning Row Identifiers
- Preventing Key Changes
- Enabling Change Capture
- Maintaining Materialized Query Tables



Triggers and Cascade Constraints Considerations

Triggers

Disable triggers on the target tables, or alter them to ignore changes made by the Oracle GoldenGate database user. Oracle GoldenGate replicates DML that results from a trigger. If the same trigger gets activated on the target table, then it becomes redundant because of the replicated version, and the database returns an error.

Cascade Constraints Considerations

Cascading updates and deletes captured by Oracle GoldenGate are not logged in binary log, so they are not captured. This is valid for both MySQL and MariaDB. For example, when you run the delete statement in the parent table with a parent child relationship between tables, the cascading deletes (if there are any) happens for child table, but they are not logged in binary log. Only the delete or update record for the parent table is logged in the binary log and captured by Oracle GoldenGate.

See https://mariadb.com/kb/en/replication-and-foreign-keys/ and https://dev.mysql.com/doc/refman/8.0/en/innodb-and-mysql-replication.html for details.

To properly handle replication of cascading operations, it is recommended to disable cascade deletes and updates on the source and code your application to explicitly delete or update the child records prior to modifying the parent record. Alternatively, you must ensure that the target parent table has the same cascade constraints configured as the source parent table, but this could lead to an out-of-sync condition between source and target, especially in cases of bi-directional replication.

Assigning Row Identifiers

Oracle GoldenGate requires some form of unique row identifier on the source and target tables to locate the correct target rows for replicated updates and deletes.

- How Oracle GoldenGate Determines the Kind of Row Identifier to Use
- Using KEYCOLS to Specify a Custom Key

How Oracle GoldenGate Determines the Kind of Row Identifier to Use

Unless a KEYCOLS clause is used in the TABLE or MAP statement, Oracle GoldenGate selects a row identifier to use in the following order of priority:

- 1. Primary key
- 2. First unique key alphanumerically that does not contain a timestamp or non-materialized computed column.
- 3. If none of the preceding key types exist (even though there might be other types of keys defined on the table) Oracle GoldenGate constructs a pseudo key of all columns that the database allows to be used in a unique key, excluding those that are not supported by Oracle GoldenGate in a key or those that are excluded from the Oracle GoldenGate configuration.



Note:

If there are other, non-usable keys on a table or if there are no keys at all on the table, Oracle GoldenGate logs an appropriate message to the report file. Constructing a key from all of the columns impedes the performance of Oracle GoldenGate on the source system. On the target, this key causes Replicat to use a larger, less efficient WHERE clause.

Using **KEYCOLS** to Specify a Custom Key

If a table does not have one of the preceding types of row identifiers, or if you prefer those identifiers not to be used, you can define a substitute key if the table has columns that always contain unique values. You define this substitute key by including a KEYCOLS clause within the Extract TABLE parameter and the Replicat MAP parameter. The specified key will override any existing primary or unique key that Oracle GoldenGate finds.

Preventing Key Changes

Do not add columns to a key after Oracle GoldenGate starts extracting data from the table. This rule applies to a primary key, a unique key, a KEYCOLS key, or an all-column key. DB2 LUW does not supply a before image for columns that are added to a table. If any columns in a key are updated on the source, Oracle GoldenGate needs a before image to compare with the current values in the target table when it replicates the update.

Enabling Change Capture

Configure DB2 to log data changes in the expanded format that is supplied by the DATA CAPTURE CHANGES feature of the CREATE TABLE and ALTER TABLE commands. This format provides Oracle GoldenGate with the entire before and after images of rows that are changed by update statements. You can use GGSCI to issue the ALTER TABLE command as follows.

To Enable Change Capture from GGSCI:

- 1. From the Oracle GoldenGate directory, run GGSCI.
- 2. Log on to DB2 from GGSCI as a user that has ALTER TABLE privileges. Specify the data source name with SOURCEDB and specify the user login with USERID and PASSWORD.

```
DBLOGIN SOURCEDB dsn, USERID user[, PASSWORD password]
```

3. Issue the following command. where <code>owner.table</code> is the fully qualified name of the table. You can use a wildcard to specify multiple table names. Only the asterisk (*) wildcard is supported for DB2 LUW.

```
ADD TRANDATA owner.table
```

ADD TRANDATA issues the following command, which includes logging the before image of LONGVAR columns:

ALTER TABLE name DATA CAPTURE CHANGES INCLUDE LONGVAR COLUMNS;

Example 2-1 To Exclude LONGVAR Logging:

To omit the INCLUDE LONGVAR COLUMNS clause from the ALTER TABLE command, use ADD TRANDATA with the EXCLUDELONG option.

ADD TRANDATA owner.table, EXCLUDELONG



If LONGVAR columns are excluded from logging, the Oracle GoldenGate features that require before images, such as the <code>GETUPDATEBEFORES</code>, <code>NOCOMPRESSUPDATES</code>, and <code>NOCOMPRESSUPLATES</code> parameters, might return errors if tables contain those columns. For a workaround, see the <code>REQUIRELONGDATACAPTURECHANGES</code> | <code>NOREQUIRELONGDATACAPTURECHANGES</code> options of the <code>TRANLOGOPTIONS</code> parameter.

Maintaining Materialized Query Tables

To maintain parity between source and target materialized query tables (MQT), you replicate the base tables, but not the MQTs. The target database maintains the MQTs based on the changes that Replicat applies to the base tables.

The following are the rules for configuring these tables:

- Include the base tables in your TABLE and MAP statements.
- Do not include MQTs in the TABLE and MAP statements.
- Wildcards can be used in TABLE and MAP statements, even though they might resolve MQT names along with regular table names. Oracle GoldenGate automatically excludes MQTs from wildcarded table lists. However, any MQT that is explicitly listed in an Extract TABLE statement by name will cause Extract to abend.

Database Configuration for DB2 LUW

- The Oracle GoldenGate Extract process calls the DB2READLOG function in the Administrative API to read the transaction log files of a DB2 LUW source database. In addition to DB2READLOG, Extract uses a small number of other API routines to check the source database configuration on startup.
- The Oracle GoldenGate Replicat process uses the DB2 CLI interface on a DB2 LUW target database. For instructions on installing this interface, see the DB2 documentation.
- The database can reside on a different server from the one where Oracle GoldenGate is installed, so long as the database is defined locally. For example, the following enables you to use database mydb locally with data that is on abc123:

```
catalog tcpip node abc123 remote abc123.us.mycompany.com server 00000 catalog db mydb as abc123 at node abc123 AUTHENTICATION server \frac{1}{2}
```

• The DB2 Universal Database has an internal trace facility called db2trc, which acquires Interprocess Communication resources (IPC) (both semaphore and shared memory). Even though a DB2 trace is not turned on, it may issue semget () calls to the operating system. These calls fail since no IPC resources are acquired so you must issue the following command on the DB2 client:

```
db2trc alloc
```

For best performance for DB2 clients with a local database, Oracle recommends that you
create a local node catalog instead of TCP/IP when connecting Oracle GoldenGate to a

database that resides on the same machine. This is because local node uses IPC, which is much faster than a TCP/IP node that uses a socket API to access the local database.

Database User for Oracle GoldenGate Processes for DB2 LUW

- Create a database user that is dedicated to Oracle GoldenGate. It can be the same user for all of the Oracle GoldenGate processes that must connect to a database:
 - Extract (source database)
 - Replicat (target database)
 - DEFGEN (source or target database)
- To preserve the security of your data, and to monitor Oracle GoldenGate processing
 accurately, do not permit other users, applications, or processes to log on as, or operate
 as, the Oracle GoldenGate database user. It is recommended that you store the login
 credentials in an Oracle GoldenGate credential store. The credential store makes use of
 local secure storage for the login names and passwords, and permits you to specify only
 an alias in the Oracle GoldenGate parameter files.
- Assign system administrator (SYSADM) or database administrator (DBADM) authority to the
 database user under which Extract runs. To give the Extract user DBADM authority, a user
 with SYSADM authority can issue the following grant statement.

```
GRANT DBADM ON DATABASE TO USER user
```

This authority can also be granted from the User and Group Objects folder in the DB2 Control Center. The database tab for the user that is assigned to an Oracle GoldenGate process should have the Database Administrative Authority box checked.

Note:

If the Extract user does not have the required authority, Extract will log the following errors and stop.

```
[SC=-1224: 3QL1224N A database agent could not be started to service a request, or was terminated as a result of a database system shurdown or a force command.

SQL STATE 55032: The CONNECT statement is invalid, because the database manager was stopped after this application was started]
```

- Grant at least the following privileges to the database user under which Replicat runs:
 - Local CONNECT to the target database
 - SELECT on the system catalog views
 - SELECT, INSERT, UPDATE, and DELETE on the target tables

Setting the Session Character Set

To support the conversion of character sets between the source and target databases, make certain that the session character set is the same as the database character set. You can set the session character set with the DB2CODEPAGE environment variable.



Preparing for Initial Extraction

During the initialization of the Oracle GoldenGate environment, you will be doing an initial data synchronization and starting the Oracle GoldenGate processes for the first time. In conjunction with those procedures, you will be creating process groups. To create an Extract group, an initial start position must be established in the transaction log. This initial read position is on a transaction boundary that is based on one of the following:

- End of the transaction file
- A specific LRI value

The start point is specified with the BEGIN option of the ADD EXTRACT command.

When the Extract process starts for the first time, it captures all the transaction data that it encounters after the specified start point, but none of the data that occurred *before* that point. This can cause partial transactions to be captured if open transactions span the start point.

To ensure initial transactional consistency:

To avoid the capture of partial transactions, initialize the Extract process at a point in time when the database is in a paused state. DB2 LUW provides a QUIESCE command for such a purpose. This is the only way to ensure transactional consistency.



After the Extract is past the initialization, subsequent restarts of the Extract do not extract partial transactions, because the process uses recovery checkpoints to mark its last read position.

To view open transactions:

IBM provides a utility called <code>db2pd</code> for monitoring DB2 databases and instances. You can use it to view information about open transactions and to determine if any of them span the start point. However, because DB2 LUW log records lack timestamps, it might not be possible to make an accurate assessment. If possible, quiesce the database prior to initialization of Oracle GoldenGate.

For more information on initializing the Oracle GoldenGate environment, see Instantiating Oracle GoldenGate with an Initial Load in *Administering Oracle GoldenGate*.

Specifying the DB2 LUW Database in Parameter Files

For an Oracle GoldenGate process to connect to the correct DB2 LUW database, you must specify the name (not an alias) of the DB2 LUW database with the following parameters:

- Specify the DB2 source database with the Extract parameter SOURCEDB.
- Specify the DB2 target database name with the Replicat parameter TARGETDB.

For more information about these parameters, see the Reference for Oracle GoldenGate for Windows and UNIX.



Configuring Oracle GoldenGate for DB2 LUW

This chapter provides an overview of the basic steps required to configure Oracle GoldenGate for a DB2 LUW source and target database.

Topics:

- What to Expect from these Instructions
- Where to Get More Information
- Configuring the Primary Extract
- Configuring the Data Pump Extract
- Creating a Temporal Table
- When to Start Replicating Transactional Changes
- Testing Your Configuration

What to Expect from these Instructions

These instructions show you how to configure basic parameter (configuration) files for the following processes:

- A primary Extract (captures transaction data from the data source)
- A data-pump Extract (propagates the data from local storage to the target system)
- A Replicat (applies replicated data to the target database)

Your business requirements probably will require a more complex topology, but this procedure forms a basis for the rest of your configuration steps.

By performing these steps, you can:

- · Get the basic configuration files established.
- Build upon them later by adding more parameters as you make decisions about features or requirements that apply to your environment.
- Use copies of them to make the creation of additional parameter files faster than starting from scratch.

Where to Get More Information

See Administering Oracle GoldenGate and Securing the Oracle GoldenGate Environment for more information about:

- The processes and files that you are configuring
- Detailed configuration information
- Security options
- Data-integration options (filtering, mapping, conversion)
- Instructions for configuring complex topologies

Steps to perform initial instantiation of the replication environment

Configuring the Primary Extract

These steps configure the primary Extract to capture transaction data from a source DB2 LUW and write the data to a local trail for temporary storage.

1. In GGSCI on the source system, create the Extract parameter file.

EDIT PARAMS name

Where: name is the name of the primary Extract.

2. Enter the Extract parameters in the order shown, starting a new line for each parameter statement.

Basic parameters for the primary Extract

EXTRACT finance
SOURCEDB mysource, USERIDALIAS myalias
ENCRYPTTRAIL AES192
EXTTRAIL /ggs/dirdat/lt
TABLE hr.*;

Parameter	Description		
EXTRACT group	group is the name of the Extract group.		
SOURCEDB database, USERIDALIAS alias	Specifies the real name of the source DB2 for i database (not an alias), plus the alias of the database login credential of the user that is assigned to Extract. This credential must exist in the Oracle GoldenGate credential store. For more information, see Database User for Oracle GoldenGate Processes.		
ENCRYPTTRAIL algorithm	Encrypts the local trail.		
EXTTRAIL pathname	Specifies the path name of the local trail to which the primary Extract writes captured data for temporary storage.		
TABLE schema.object;	 Specifies the database object for which to capture data. TABLE is a required keyword. schema is the schema name or a wildcarded set of schemas. object is the table name, or a wildcarded set of tables. The question mark (?) wildcard is not supported for this database. Note that only the asterisk (*) wildcard is supported for DB2 LUW. Terminate the parameter statement with a semi-colon. To exclude tables from a wildcard specification, use the TABLEEXCLUDE parameter. For more information and for additional options that control data filtering, mapping, and manipulation, see TABLE MAP in Reference for Oracle GoldenGate. 		

- Enter any optional Extract parameters that are recommended for your configuration. You
 can edit this file at any point before starting processing by using the EDIT PARAMS
 command in GGSCI.
- 4. Save and close the file.



Configuring the Data Pump Extract

These steps configure the data pump that reads the local trail and sends the data across the network to a remote trail on the target. The data pump is optional, but recommended.

1. In GGSCI on the source system, create the data-pump parameter file.

```
EDIT PARAMS name
```

Where name is the name of the data-pump Extract.

 Enter the data-pump Extract parameters in the order shown, starting a new line for each parameter statement. Your input variables will be different. See #unique_47/ unique 47 Connect 42 CEGHIGIC for descriptions.

Basic parameters for the data-pump Extract group:

```
EXTRACT extpump

SOURCEDB mypump, USERIDALIAS myalias

RMTHOST fin1, MGRPORT 7809 ENCRYPT AES192, KEYNAME securekey2

RMTTRAIL /ggs/dirdat/rt

TABLE hr.*;
```

Parameter	Description				
EXTRACT group	group is the name of the data pump Extract.				
SOURCEDB database, USERIDALIAS alias	Specifies the real name of the source DB2 LUW database (not an alias), plus the alias of the database login credential of the user that is assigned to Extract. This credential must exist in the Oracle GoldenGate credential store. For more information, see Database User for Oracle GoldenGate Processes.				
RMTHOST hostname, MGRPORT portnumber, [, ENCRYPT algorithm KEYNAME keyname]	 RMTHOST specifies the name or IP address of the target system. MGRPORT specifies the port number where Manager is running on the target. ENCRYPT specifies optional encryption of data across TCP/IP. 				
RMTTRAIL pathname	Specifies the path name of the remote trail.				
TABLE schema.object;	Specifies a table or sequence, or multiple objects specified with a wildcard. In most cases, this listing will be the same as that in the primary Extract parameter file. Table is a required keyword. schema is the schema name or a wildcarded set of schemas. object is the name of a table or a wildcarded set of tables. Only the asterisk (*) wildcard is supported for DB2 LUW. The question mark (?) wildcard is not supported for this database. See Oracle Golden Gate Administering Oracle GoldenGate for Windows and UNIX for information about how to specify object names with and without wildcards. Terminate the parameter statement with a semi-colon. To exclude tables from a wildcard specification, use the TableEXCLUDE parameter. For more information and for additional Table options that control data filtering, mapping, and manipulation, see Oracle Fusion Middleware Reference for Oracle GoldenGate for Windows and UNIX.				



- Enter any optional Extract parameters that are recommended for your configuration. You
 can edit this file at any point before starting processing by using the EDIT PARAMS
 command in GGSCI.
- 4. Save and close the file.

Creating a Temporal Table

A temporal table is a table that maintains the history of its data and the time period when its data are valid. Temporal tables are used in Oracle GoldenGate to keep track of all the old rows that are deleted or updated in the table. Temporal tables are also used to maintain the business validity of its rows and data. For example, Oracle GoldenGate keeps track of the time period during which a row is valid. There are three types of temporal tables, system-period, application-period, and bitemporal table.

- Support for Temporal Tables
- · Replicating with Temporal Tables
- Converting
- Creating a Checkpoint Table
- Configuring the Replicat Parameter File

Support for Temporal Tables

- Replication between system-period temporal tables and application-period temporal tables is not supported.
- Replication from a non-temporal table to a temporal table is not supported.
- Replication of temporal tables with the INSERTALLRECORDS parameter is not supported.
- Bidirectional replication is supported only with the default replication.
- CDR in bidirectional replication is not supported.
- CDR in application-period temporal tables is supported.

Replicating with Temporal Tables

You can choose one of the following methods to replicate a system-period or a bitemporal temporal table as follows:

- You can replicate a temporal table to another temporal table only; this is the default behavior. Oracle GoldenGate will not replicate the SYSTEM_TIME period and transaction id columns because these are automatically generated columns at the apply side. The database manager populates the columns in the target temporal table using the system clock time and with the default values. You can preserve the original values these columns then use any of the following:
 - Add extra timestamp columns in the target temporal table and map the columns accordingly. The extra columns are automatically added to the associated history table.
 - Use a non-temporal table at the apply side and map the columns appropriately. In this scenario, you will not be able to maintain the history table.
 - In a heterogeneous configuration where the source is DB2 LUW and the target is a different database, you can either ignore the automatically generated columns or use



an appropriate column conversion function to convert the columns value in the format that target database supports and map them to target columns accordingly.

Or

You can replicate a temporal table, with the associated history table, to a temporal and history table respectively then you must specify the replicate parameter, DBOPTIONS SUPPRESSTEMPORALUPDATES. You must specify both the temporal table and history table to be captured in the Extract parameter file. Oracle GoldenGate replicates the SYSTEM_TIME period and transactions id columns value. You must ensure that the database instance has the execute permission to run the stored procedure at the apply side.

Oracle GoldenGate cannot detect and resolve conflicts while using default replication as SYSTEM_TIME period and transactionstart id columns remains auto generated. These columns cannot be specified in set and where clause. If you use the SUPPRESSTEMPORALUPDATES parameter, then Oracle GoldenGate supports CDR.

Converting

You can convert an already existing table into a temporal table, which changes the structure of the table. This section describes how the structure of the tables changes. The following sample existing table is converted into all three temporal tables types in the examples in this section:

Example 1 Converting an existing table into System-period temporal table.

You convert the sample existing table into a system-period temporal table by adding SYSTEM_PERIOD, transaction id columns, and SYSTEM_TIME period as in the following:

```
ALTER TABLE policy_info
ADD COLUMN sys_start TIMESTAMP(12) NOT NULL GENERATED ALWAYS AS ROW BEGIN;
ALTER TABLE policy_info
ADD COLUMN sys_end TIMESTAMP(12) NOT NULL GENERATED ALWAYS AS ROW END;
ALTER TABLE policy_info
ADD COLUMN ts_id TIMESTAMP(12) NOT NULL GENERATED ALWAYS AS TRANSACTION START ID;
ALTER TABLE policy_info ADD PERIOD SYSTEM_TIME(sys_start, sys_end);
```

Then you create a history table for the new temporal table using one of the following two methods:



CREATE TABLE hist policy info LIKE policy info with RESTRICT ON DROP;

The RESTRICT ON DROP clause will not allow the history table to get dropped while dropping system-period temporal table. Otherwise the history table gets implicitly dropped while dropping its associated temporal table. You can create a history table without RESTRICT ON DROP. A history table cannot be explicitly dropped.

You should not use the <code>GENERATED ALWAYS</code> clause while creating a history table. The primary key of the system-period temporal table also does not apply here as there could be many updates for a particular row in the base table, which triggers many inserts into the history table for the same set of primary keys. Apart from these, the structure of a history table should be exactly same as its associated system-period temporal table. The history table must have the same number and order of columns as system-period temporal table. History table columns cannot explicitly be added, dropped, or changed. You must associate a system-period temporal table with its history table with the following statement:

```
ALTER TABLE policy_info ADD VERSIONING USE HISTORY TABLE hist_policy_info.
```

The GENERATED ALWAYS columns of the table are the ones that are always populated by the database manager so you do not have any control over these columns. The database manager populates these columns based on the system time.

The extra added SYSTEM_PERIOD and transaction id columns will have default values for already existing rows as in the following:

The associated history table is populated with the before images once you start updating the temporal table.

Example 2 Converting an existing table into application-period temporal table.

You can convert the sample existing table into application-period temporal table by adding time columns and a BUSINESS TIME period as in the following:

```
ALTER TABLE policy_info ADD COLUMN bus_start DATE NOT NULL DEFAULT '10/10/2001'" ALTER TABLE policy_info ADD COLUMN bus_end DATE NOT NULL DEFAULT '10/10/2002' ALTER TABLE policy_info ADD PERIOD BUSINESS_TIME(bus_start, bus_end)
```

While adding time columns, you need to make sure that while entering business validity time values of the existing time columns, the <code>bus_start</code> column always has value lesser than <code>bus_end</code> because these columns specify the business validity of the rows.

The new application-period temporal table will look similar to:



POLICY_ID	COVERAGE	BUS_START	BUS_END	
ERT	14000		10/10/2001	10/10/2002
DEF	13000		10/10/2001	10/10/2002
ABC	12000		10/10/2001	10/10/2002

Example 3 Converting an existing table into bitemporal table.

You can convert the sample existing table into bitemporal table by adding System_period, time columns along with the system time and business time period as in the following:

```
ALTER TABLE policy_info
   ADD COLUMN sys_start TIMESTAMP(12) NOT NULL GENERATED ALWAYS AS ROW BEGIN;
ALTER TABLE policy_info
   ADD COLUMN sys_end TIMESTAMP(12) NOT NULL GENERATED ALWAYS AS ROW END;
ALTER TABLE policy_info
   ADD COLUMN ts_id TIMESTAMP(12) NOT NULL GENERATED ALWAYS AS TRANSACTION START ID;
ALTER TABLE policy_info ADD PERIOD SYSTEM_TIME(sys_start, sys_end);

ALTER TABLE policy_info ADD COLUMN bus_start DATE NOT NULL DEFAULT '10/10/2001'"
ALTER TABLE policy_info ADD COLUMN bus_end DATE NOT NULL DEFAULT '10/10/2002'
ALTER TABLE policy info ADD PERIOD BUSINESS TIME(bus start, bus end)
```

While adding the time columns, you must make sure that while entering business validity time values of already existing time columns, the <code>bus_start</code> column always has value lesser than <code>bus_end</code> because these columns specify the business validity of the rows.

Then you create a history table for the new temporal table using one of the following two methods:

• The RESTRICT ON DROP clause will not allow the history table to get dropped while dropping system-period temporal table. Otherwise the history table gets implicitly dropped while dropping its associated temporal table. You can create a history table without RESTRICT ON DROP. A history table cannot be explicitly dropped.

You should not use the GENERATED ALWAYS clause while creating a history table. The primary key of the system-period temporal table also does not apply here as there could be many updates for a particular row in the base table, which triggers many inserts into the history table for the same set of primary keys. Apart from these, the structure of a history table should be exactly same as its associated system-period temporal table. The history table must have the same number and order of columns as system-period temporal table. History table columns cannot explicitly be added, dropped, or changed. You must associate a system-period temporal table with its history table with the following statement:

ALTER TABLE policy info ADD VERSIONING USE HISTORY TABLE hist policy info.



The GENERATED ALWAYS columns of the table are the ones that are always populated by the database manager so you do not have any control over these columns. The database manager populates these columns based on the system time.

The extra added SYSTEM_PERIOD and transaction id columns will have default values for already existing rows as in the following:

The associated history table is populated with the before images once you start updating the temporal table.

The extra added SYSTEM_TIME period, transaction id and time columns will have default values for already existing rows as in the following:

The history table is populated with the before images once user starts updating the temporal table.

Example 4 Replication in Heterogeneous Environment.

In heterogeneous configuration in which you do not have temporal tables at the apply side, you can only replicate the system-period and bitemporal tables though *not* the associated history tables. While performing replication in this situation, you must take care of the SYSTEM_PERIOD and transaction id columns value. These columns will have some values that the target database might not support. You should first use the map conversion functions to convert these values into the format that the target database supports, and then map the columns accordingly.



To replicate the row into MySQL, you would use the colmap() function:

```
map source_schema.policy_info, target target_schema.policy_info colmap
(policy_id=policy_id, coverage=coverage, sys_start= @IF( ( @NUMSTR( @STREXT(sys_start,1,4))) > 1000, sys_start, '1000-01-01 00.00.00.000000'), sys_end=sys_end,
    ts_id= @IF( ( @NUMSTR( @STREXT(ts_id,1,4))) > 1000, ts_id, '1000-01-01
    00.00.00.000000'));
```

Creating a Checkpoint Table

The checkpoint table is a required component of Replicat.

Replicat maintains its recovery checkpoints in the checkpoint table, which is stored in the target database. Checkpoints are written to the checkpoint table within the Replicat transaction. Because a checkpoint either succeeds or fails with the transaction, Replicat ensures that a transaction is only applied once, even if there is a failure of the process or the database.

To configure a checkpoint table, see Creating a Checkpoint Table in *Administering Oracle GoldenGate*.

Configuring the Replicat Parameter File

These steps configure the Replicat process. This process applies replicated data to a DB2 LUW target database.

1. In GGSCI on the target system, create the Replicat parameter file.

```
EDIT PARAMS name
```

Where: name is the name of the Replicat group.

Enter the Replicat parameters in the order shown, starting a new line for each parameter statement.

Basic parameters for the Replicat group:

```
REPLICAT financer
TARGETDB mytarget, USERIDALIAS myalias
ASSUMETARGETDEFS
MAP hr.*, TARGET hr2.*;
```

Parameter	Description
REPLICAT group	group is the name of the Replicat group.
TARGETDB database, USERIDALIAS alias	Specifies the real name of the target DB2 LUW database (not an alias), plus the alias of the database login credential of the user that is assigned to Replicat. This credential must exist in the Oracle GoldenGate credential store. For more information, see Database User for Oracle GoldenGate Processes.



Parameter	Description
ASSUMETARGETDEFS	Specifies how to interpret data definitions. ASSUMETARGETDEFS assumes the source and target tables have identical definitions. (This procedure assume identical definitions.)
	Use the alternative SOURCEDEFS if the source and target tables have different definitions, and create a source data-definitions file with the DEFGEN utility.
MAP schema.object, TARGET schema.object;	 Specifies the relationship between a source table or multiple objects, and the corresponding target object or objects. MAP specifies the source portion of the MAP statement and is a required keyword. Specify the source objects in this clause. TARGET specifies the target portion of the MAP statement and is a required keyword. Specify the target objects to which you are mapping the source objects. schema is the schema name or a wildcarded set of schemas. object is the name of a table or a wildcarded set of objects. Terminate this parameter statement with a semi-colon. Note that only the asterisk (*) wildcard is supported for DB2 LUW. The question mark (?) wildcard is not supported for this database. To exclude objects from a wildcard specification, use the MAPEXCLUDE parameter.

- 3. Enter any optional Replicat parameters that are recommended for your configuration. You can edit this file at any point before starting processing by using the EDIT PARAMS command in GGSCI.
- Save and close the file.

When to Start Replicating Transactional Changes

You must start replication when the source and target data is in a synchronized state, where the corresponding rows in the source and target tables contain identical data values. Unless you are starting with brand new source and target databases with no current user activity, you will need to activate change capture and apply processes to handle ongoing transactional changes while an initial load is being applied to the target. This process is known as *initial synchronization*, or also as *instantiation*. The initial load captures a point-in-time snapshot of the source data and applies it to the target, while Oracle GoldenGate maintains any changes that are made after that point.

See Instantiating Oracle GoldenGate with an Initial Load in *Administering Oracle GoldenGate* for instantiation options.

Testing Your Configuration

It is important to test your configuration in a test environment before deploying it live on your production machines. This is especially important in an active-active or high availability configuration, where trusted source data may be touched by the replication processes. Testing enables you to find and resolve any configuration mistakes or data issues without the need to interrupt user activity for re-loads on the target or other troubleshooting activities.



Part III

Using Oracle GoldenGate for DB2 for i

With Oracle GoldenGate for DB2 for i, you can perform initial loads and capture transactional data from supported DB2 for i versions and replicate the data to a DB2 for i database or other supported Oracle GoldenGate targets, such as an Oracle Database.

Oracle GoldenGate for DB2 for i supports data filtering, mapping, and transformations unless noted otherwise in this documentation. Perform initial loads from DB2 for i to target tables in DB2 for i or other databases to instantiate a synchronized replication environment.

This part describes tasks for configuring and running Oracle GoldenGate for DB2 for i.

Topics:

- Understanding What's Supported for DB2 for i
 This chapter contains information on database and table features supported by Oracle GoldenGate for DB2 for i.
- Preparing the System for Oracle GoldenGate
- · Configuring Oracle GoldenGate for DB2 for i
- Instantiating and Starting Oracle GoldenGate Replication
- Using Remote Journal



4

Understanding What's Supported for DB2 for i

This chapter contains information on database and table features supported by Oracle GoldenGate for DB2 for i.

Oracle GoldenGate on DB2 for i supports the filtering, mapping, and transformation of data unless otherwise noted in this documentation.

Oracle GoldenGate for DB2 for i runs directly on a DB2 for i source system to capture data from the transaction journals for replication to a target system. To apply data to a target DB2 for i database, Oracle GoldenGate can run directly on the DB2 for i target system or on a remote Windows or Linux system. If installed on a remote system, Replicat delivers the data by means of an ODBC connection, and no Oracle GoldenGate software is installed on the DB2 for i target.



The DB2 for i platform uses one or more **journals** to keep a record of transaction change data. For consistency of terminology in the supporting administrative and reference Oracle GoldenGate documentation, the terms "log" or "transaction log" may be used interchangeably with the term "journal" where the use of the term "journal" is not explicitly required.

Topics:

- Supported DB2 for i Data Types
- Non-Supported DB2 for i Data Types
- Supported Objects and Operations for DB2 for i
- Non-Supported Objects and Operations for DB2 for i
- Oracle GoldenGate Parameters Not Supported for DB2 for i
- Supported Object Naming Conventions
- System Schemas for DB2 for i
- Supported Character Sets

Supported DB2 for i Data Types

Oracle GoldenGate supports all DB2 for i data types, except those listed in Non-Supported DB2 for i Data Types.

Limitations of support

Extract fully supports the capture and apply of <code>TIMESTAMP(0)</code> through <code>TIMESTAMP(12)</code> when the output trail format is 19.1 or higher. Otherwise Extract truncates data from <code>TIMESTAMP(10)</code> through <code>TIMESTAMP(12)</code> to nanoseconds (maximum of nine digits of fractional time) and issues a warning to the report file.

Oracle GoldenGate supports timestamp data from 0001/01/03:00:00:00:00.000000 to 9999/12/31:23:59:59.999999. If a timestamp is converted from GMT to local time, these limits also apply to the resulting timestamp. Depending on the time zone, conversion may add or subtract hours, which can cause the timestamp to exceed the lower or upper supported limit.

Non-Supported DB2 for i Data Types

Oracle GoldenGate does not support the following DB2 for i data types:

- XML
- DATALINK
- User-defined types

Supported Objects and Operations for DB2 for i

Oracle GoldenGate supports the following DB2 for i objects and operations.

- Extraction and replication of DML operations .
- Tables with the maximum row length supported by the database.
- Tables that contain up to the maximum number of columns that is supported by the database, up to the maximum supported column size.
- DELETE FROM with no WHERE clause SQL statements and Clear Physical File Member (CLRPFM)
- Base tables underlying Materialized Query Tables, but not the MQTs themselves. The
 target database automatically maintains the content of the MQT based on the changes that
 are applied to the base tables by Replicat.
- Both Library (Native) names including members, and SQL names are allowed.
- Partitioned tables
- Supported options with SHOWTRANS:

```
SHOWTRANS [transaction_ID] [COUNT n]
[DURATION duration unit] [TABULAR]
[FILE file name] |
```

Options for SKIPTRANS and FORCETRANS:

```
SKIPTRANS transaction_ID [FORCE]
FORCETRANS transaction ID [FORCE]
```

- Limitations on Automatic Heartbeat Table support are as follows:
 - The ADD HEARTBEATTABLE command creates a new file called ogghbfreq; in the Oracle GoldenGate build folder. Do not delete this file because the pase heartbeat program reads the frequency values from it.
 - There is an extra executable in the Oracle GoldenGate build folder named ogghb.
 - An extra process named ogghb starts running from the time the ADD HEARTBEATTABLE command is given and runs until you disable the heartbeat with the DELETE HEARTBEATTABLE command. This process automatically restarts even if it is killed. To remove this process from the system, use the DELETE HEARTBEATTABLE command.



- When using the ALTER HEARTBEATTABLE command to change the heartbeat frequency with the PURGE_FREQUENCY or RETENTION_TIME options, it takes approximately 60 + older 'frequency') seconds to be implemented.
- There is an initial delay of 30 seconds between ADD HEARTBEATTABLE and the first record is updated in the heartbeat seed table.
- [THREAD n] and [DETAIL] is not supported.

Non-Supported Objects and Operations for DB2 for i

Oracle GoldenGate does not support the following objects or operations for DB2 for i.

- DDL operations
- Schema, table or column names that have trailing spaces.
- Multiple instances of a database
- The Multi-Journal feature does not support multi-journal sync of a transaction across multiple journals.

Oracle GoldenGate Parameters Not Supported for DB2 for i

This section lists some of the Oracle GoldenGate configuration parameters that are not supported for the DB2 for i platform. For full descriptions of Oracle GoldenGate parameters and the databases they support, see Oracle GoldenGate Parameters.

```
BATCHSQL (not supported on V5R4 and i6.1 only)
BR
ASCIITOEBCDIC and EBCDICTOASCII
BINARYCHARS
LOBMEMORY
TRAILCHARSETEBCDIC
Any encryption options that use AES encryption
```

Supported Object Naming Conventions

Oracle GoldenGate supports SQL naming conventions and also supports native file system names in the format of <code>library/file(member)</code>.

- library/*all(*all)
- library/a*(a*)
- library/abcde*

Oracle GoldenGate does not allow wildcarding for library names.

The member name is optional and may be left off. In that case, data for all of the members will be extracted, but only the library and file names will be captured and included in the records that are written to the trail. The result is that the data will appear to have come from only one member on the source, and you should be aware that this could cause integrity conflicts on the



target if there are duplicate keys across members. To include the member name in the trail records, include the member explicitly or though a wildcarded member specification.

For SQL names, only the first member in the underlying native file is extracted in accordance with the normal operation of SQL on an DB2 for i system. For SQL names, Oracle GoldenGate supports the wildcarding of table names, but not schema names. For instructions on wildcarding SQL names, see Specifying Object Names in Oracle GoldenGate Input in *Administering Oracle GoldenGate*.

System Schemas for DB2 for i

The following schemas or objects are not automatically replicated by Oracle GoldenGate unless they are explicitly specified without a wildcard..

- "0*"
- "SYSIBM"
- "SYSIBMADM"
- "SYSPROC"
- "SYSTOOLS"
- "#LIBRARY"
- "#RPGLIB"

Supported Character Sets

The default behavior of a DB2 for i Extract is to convert all character data to Unicode. The overhead of the performance of the conversion to UTF-8 for the text data has been substantially reduced. However, if you want to send data in its native character set you can use the CHARSET and COLCHARSET parameters to override the default behavior as follows. See CHARSET and COLCHARSET in *Reference for Oracle GoldenGate*.



Preparing the System for Oracle GoldenGate

This chapter contains guidelines for preparing the DB2 for i system to support Oracle GoldenGate.

Topics:

- Preparing the Journals for Data Capture by Extract
- Specifying Object Names
- Preparing Tables for Processing
- Adjusting the System Clock
- · Configuring the ODBC Driver

Preparing the Journals for Data Capture by Extract

All tables for which you want data to be captured must be journaled, either explicitly or by default by means of a QSQJRN journal in the same library. To preserve data integrity, data journal entries are sent to the Extract process in time order as they appear on the system. This section provides guidelines for configuring the journals to support capture by the Extract process.

- Allocating Journals to an Extract Group
- Setting Journal Parameters
- Deleting Old Journal Receivers

Allocating Journals to an Extract Group

One Extract process can handle up to 30 journals. If using more journals than that, use additional Extract processes to handle the extra journals. You can also use additional Extract processes to improve capture performance if necessary.



To ensure transaction integrity, all journals that correspond to any given transaction must be read by the same Extract group. For more information about using multiple Extract processes, see Tuning the Performance of Oracle GoldenGate in *Administering Oracle GoldenGate*.

Setting Journal Parameters

To support the capture of data by the Extract process, the following are the minimal journaling parameter settings that are required.

Manage Receivers (MNGRCV): *SYSTEM

- Delete Receivers (DLTRCV): *NO
- Receiver Size Option (RCVSIZOPT): *MAXOPT2 (*MAXOPT3 recommended)
- Journal State (JRNSTATE): *ACTIVE
- Minimize Entry Specific Data (MINENTDTA): *NONE
- Fixed Length Data (FIXLENDTA): *USR *SYSSEQ

In the following example, the command to set these attributes for a journal $\tt JRN1$ in library LIB1 would be:

CHGJRN JRN(LIB1/JRN1) MNGRCV(*SYSTEM) DLTRCV(*NO) RCVSIZOPT(*MAXOPT2) JRNSTATE(*ACTIVE) MINENTDTA(*NONE) FIXLENDTA(*USR *SYSSEQ)



To check the attributes of a journal, use the command $WRKJRNA\ JRN\ (LIB1/JRN1)\ DETAIL\ (*CURATR)$.

When the journaling is set to the recommended parameter settings, you are assured that the entries in the journals contain all of the information necessary for Oracle GoldenGate processing to occur. These settings also ensure that the system does not delete the journal receivers automatically, but retains them in case Extract needs to process older data.

Deleting Old Journal Receivers

Although the DLTRCV parameter is set to NO in the recommended configuration for Extract (see Setting Journal Parameters), you can delete old journal receivers manually once Extract is finished capturing from them.

If using another application that is using the journals that Oracle GoldenGate will be reading, consideration must be given regarding any automatic journal receiver cleanup that may be in place. Oracle GoldenGate must be able to read the journal receivers before they are detached or removed.

To Delete Journal Receivers

- Run GGSCI.
- 2. In GGSCI, issue the following command to view the journal positions in which Extract has current processing points, along with their journal receivers.

```
INFO EXTRACT group
```

3. Use the following DB2 for i command to delete any journal receivers that were generated prior to the ones that are shown in the INFO EXTRACT command.

```
DLTJRNRCV JRNRCV(library/journal receiver)
```

Where:

library and journal_receiver are the actual names of the library and journal receiver to be deleted. See the DB2 for i Information Center for more information about this command.



Specifying Object Names

Oracle GoldenGate commands and parameters support input in the form of SQL names, native names in the format of <code>library_name/file_name(member_name)</code>, or a mix of the two. If a native file system name does not include the member name, all members are implicitly selected by the Oracle GoldenGate process. For a SQL name only the first member is used.

To support case sensitivity of double quoted object names, specify those names within double quotes in the Oracle GoldenGate parameter files. This is true of both SQL and native file system names.

When specifying a native table name in a MAP statement on a platform other than DB2 for i, the name must be enclosed within double quotes so that Oracle GoldenGate correctly interprets it as a separator character.

For consistency of terminology in other administrative and reference Oracle GoldenGate documentation, the SQL terms "schema" and "table" are used to reference the containers for the DB2 for i data, as shown here.

Table 5-1 Native-SQL object name relationships

Native	SQL	Notes
Library (maximum length 10)	Schema (maximum length 128)	The operating system creates a corresponding native name for a SQL-created schema.
File (maximum length 10)	Table (maximum length 128)	The operating system creates a corresponding native name for a SQL-created table.
Member	Not Applicable	Contains the actual data. Only the first member of a FILE object can be accessed through SQL. To access data in other members the native system name must be used.

Preparing Tables for Processing

The following table attributes must be addressed in an Oracle GoldenGate environment.

- Assigning Row Identifiers
- · Preventing Key Changes
- Disabling Constraints on the Target
- Enabling Change Capture
- Maintaining Materialized Query Tables
- Specifying the Oracle GoldenGate Library

Assigning Row Identifiers

Oracle GoldenGate requires some form of unique row identifier on the source and target tables to locate the correct target rows for replicated updates and deletes.



- How Oracle GoldenGate Determines the Kind of Row Identifier to Use
- Using KEYCOLS to Specify a Custom Key

How Oracle GoldenGate Determines the Kind of Row Identifier to Use

Unless a KEYCOLS clause is used in the TABLE or MAP statement, Oracle GoldenGate selects a row identifier to use in the following order of priority:

- Primary key
- First unique key alphanumerically that does not contain a timestamp or non-materialized computed column.
- 3. If none of the preceding key types exist (even though there might be other types of keys defined on the table) Oracle GoldenGate constructs a pseudo key of all columns that the database allows to be used in a unique key, excluding those that are not supported by Oracle GoldenGate in a key or those that are excluded from the Oracle GoldenGate configuration.

Note:

If there are other, non-usable keys on a table or if there are no keys at all on the table, Oracle GoldenGate logs an appropriate message to the report file. Constructing a key from all of the columns impedes the performance of Oracle GoldenGate on the source system. On the target, this key causes Replicat to use a larger, less efficient WHERE clause.

Using **KEYCOLS** to Specify a Custom Key

If a table does not have one of the preceding types of row identifiers, or if you prefer those identifiers not to be used, you can define a substitute key if the table has columns that always contain unique values. You define this substitute key by including a KEYCOLS clause within the Extract TABLE parameter and the Replicat MAP parameter. The specified key will override any existing primary or unique key that Oracle GoldenGate finds.

Preventing Key Changes

If you must add columns to the key that Extract is using as the row identifier for a table (primary key, unique key, KEYCOLS key, or all-column key) after Oracle GoldenGate has started processing journal data, follow these steps to make the change.

Stop Extract.

STOP EXTRACT group

Issue the following command until it returns EOF, indicating that it has processed all of the existing journal data.

INFO EXTRACT group

- 3. Make the change to the key.
- 4. Start Extract.

START EXTRACT group



Disabling Constraints on the Target

Triggers and cascade constraints must be disabled on the target tables or configured to ignore changes made by Replicat. Constraints must be disabled because Oracle GoldenGate replicates DML that results from a trigger or a cascade constraint. If the same trigger or constraint gets activated on the target table, it becomes redundant because of the replicated version, and the database returns an error. Consider the following example, where the source tables are emp_src and salary_src and the target tables are emp_targ and salary_targ.

- A delete is issued for emp src.
- 2. It cascades a delete to salary src.
- 3. Oracle GoldenGate sends both deletes to the target.
- The parent delete arrives first and is applied to emp targ.
- 5. The parent delete cascades a delete to salary targ.
- **6.** The cascaded delete from salary_src is applied to salary_targ.
- The row cannot be located because it was already deleted in step 5.

Enabling Change Capture

To capture changes to a table in a journal, you can run the STRJRNPF command on the OS/400 command line or run the ADD TRANDATA command from GGSCI. The ADD TRANDATA command calls STRJRNPF and is the recommended method to start journaling for tables, because it ensures that the required journal image attribute of Record Images (IMAGES): *BOTH is set on the STRJRNPF command.

To Run add trandata

- Run GGSCI on the source system.
- Issue the DBLOGIN command.

DBLOGIN SOURCEDB database USERID user, PASSWORD password [encryption options]

Where: SOURCEDB specifies the default DB 2 for i database, USERID specifies the Extract user profile, and PASSWORD specifies that profile's password.



Only BLOWFISH encryption is supported for DB2 for i systems.

3. Issue the ADD TRANDATA command.

ADD TRANDATA table specification

Where: table specification is one of the following:

- schema.table [JOURNAL library/journal]
- library/file [JOURNAL library/journal] (See Specifying a Default Journal)
- Specifying a Default Journal



Removing a Default Journal Specification

Specifying a Default Journal

To specify a default journal for multiple tables or files in the ADD TRANDATA command, instead of specifying the JOURNAL keyword, use the following GGSCI command before issuing ADD TRANDATA.

DEFAULTJOURNAL library/journal

Any ADD TRANDATA command used without a journal assumes the journal from DEFAULTJOURNAL.

To display the current setting of DEFAULTJOURNAL, you can issue the command with no arguments.

Removing a Default Journal Specification

To remove the use of a default journal, use the following GGSCI command:

DEFAULTJOURNAL CLEAR

Maintaining Materialized Query Tables

To maintain parity between source and target materialized query tables (MQT), you replicate the base tables, but not the MQTs. The target database maintains the MQTs based on the changes that Replicat applies to the base tables.

The following are the rules for configuring these tables:

- Include the base tables in your TABLE and MAP statements.
- Do not include MQTs in the TABLE and MAP statements.
- Wildcards can be used in TABLE and MAP statements, even though they might resolve MQT names along with regular table names. Oracle GoldenGate automatically excludes MQTs from wildcarded table lists. However, any MQT that is explicitly listed in an Extract TABLE statement by name will cause Extract to abend.

Specifying the Oracle GoldenGate Library

Before starting Oracle GoldenGate, specify the name of the Oracle GoldenGate for DB2 for i library when running the ggos400install script. This creates a link to the OGGPRCJRN *SRVPGM (service program) object that was restored to that library. If the link to the OGGPRCJRN *SRVPGM program is deleted you could just re-run the ggos400install shell script and specify the same library, or use the command "ln -s /qsys.lib/ $OGG_library.lib/oggprcjrn.srvpgm$ Oggprcjrn.srvpgm". If this link is incorrect or missing, Extract will abend.

Adjusting the System Clock

It is recommended that you set the system clock to UTC (Universal Time Coordinate) time and use the timezone offset in the DB2 for i system values to represent the correct local time. If this setup is done correctly, local daylight savings time adjustments can occur automatically with no disruption to replication.



Configuring the ODBC Driver

This section applies only if you installed Replicat on a Windows or Linux system to operate remotely from the DB2 for i target. In this configuration, Replicat must connect to the target system over a database connection that is specified with ODBC. The following steps show how to install and configure ODBC to connect to the DB2 for i target system.

Configuring ODBC on Linux

Configuring ODBC on Windows

- Configuring ODBC on Linux
- Configuring ODBC on Windows

Configuring ODBC on Linux

To configure ODBC, you can use the graphical user interface to run the ODBC configuration utility that is supplied with your Linux distribution, or you can edit the odbc.ini file with the settings described in these steps. (These steps show the ODBC Administrator tool launched from the ODBCConfig graphical interface utility for Linux.)

- Download and install the 32-bit or 64-bit iSeries Access ODBC driver on the remote Linux system according to the vendor documentation. The iSeries ODBC driver is supplied as a free component of iSeries Access.
- 2. Issue one of the following commands, depending on the driver that you want to use.

32-bit driver:

```
rpm -ivh iSeriesAccess-7.1.0-1.0.i386.rpm
```

64-bit driver:

```
rpm -ivh iSeriesAccess-7.1.0-1.0.x86 64.rpm
```

- 3. You can create a user DSN (a connection that is available only to the user that created it) or a system DSN (a connection that is available to all users on the system). To create a user DSN, log on to the system as the user that you will be using for the Replicat process.
- 4. Run the ODBC configuration utility.
- On the initial page of the ODBC configuration tool, select the User DSN tab to create a
 user DSN or the System DSN tab to create a system DSN. (These steps create a user
 DSN; creating a system DSN is similar.)
- On the tab you selected, click Add.
- Select the appropriate iSeries Access ODBC driver, click OK. If the correct driver is not shown in the Select the DRIVER list, click the Add button and then complete the fields.

Steps to complete the Driver Properties fields when the driver is not found:

- Set Name to the name of the driver.
- Set **Driver** to the path where the driver is installed.
- Set Setup to the libewbodbei.so file that is in the driver installation directory.
- Leave the other settings to their defaults.
- Click the check mark above the Name field to save your settings.



- 8. In the Name field of the Data Source Properties dialog, supply a one-word name for the data source. In the System field, enter the fully qualified name of the target DB2 for i system, for example: sysname.company.com. Leave the UserID and Password fields blank, to allow Replicat to supply credentials when it connects to the database. Leave the remaining fields set to their defaults, and then click the check mark above the Name field to save your settings.
- You are returned to the ODBC Data Source Administrator dialog. Click OK to exit the ODBC configuration utility.
- 10. To support GRAPHIC, VARGRAPHIC and DBCLOB types, edit the .odbc.ini file and add the following line.

GRAPHIC = 1



If you created a user Data Source Name, this file is located in the home directory of the user that created it. If you created a system DSN, this file is in /etc/odbc.ini or /usr/local/etc/odbc.ini.

11. From the Oracle GoldenGate directory on the target, run GGSCI and issue the DBLOGIN command to log into the target database.

DBLOGIN SOURCEDB database, USERID db user [, PASSWORD pw [encryption options]]

Where:

- SOURCEDB database specifies the new Data Source Name.
- USERID db user, PASSWORD pw are the Replicat database user profile and password.
- encryption options is optional password encryption.



Only BLOWFISH encryption is supported for DB2 for i systems.

Configuring ODBC on Windows

On Windows, the ODBC Administration tool is in the Administrative Tools folder as **Data Sources (ODBC)**.

- Download and install the 32-bit or 64-bit iSeries Access ODBC driver from the DB2 for iSeries Access package on the remote Windows system according to the vendor documentation. The iSeries ODBC driver is supplied as a free component of iSeries Access.
- 2. You can create a user DSN (a connection that is available only to the user that created it) or a system DSN (a connection that is available to all users on the system). To create a user DSN, log on to the system as the user that you will be using for the Replicat process.
- 3. From the Windows Control Panel, select **Administrative Tools**, then **Data Sources** (ODBC).



- 4. On the first page of the ODBC configuration tool, select the User DSN tab to create a user DSN or the System DSN tab to create a system DSN. (These steps create a user DSN; creating a system DSN is similar.)
- 5. On the tab that you selected, click **Add**.
- Select the appropriate iSeries Access ODBC Driver from the list of drivers, and then click Finish.
- 7. On the General tab of the DB2 for i Access for Windows ODBC Setup dialog, provide a name (without any spaces) in the **Data Source Name** field, add an optional description in the **Description** field, and then select the system name from the **System** selection list.
- 8. On the Server tab, set Naming Convention to SQL Naming Convention (*SQL). Leave the other fields set to their defaults.
- On the Data Types tab, select the Report as Supported check box under Double Byte Character Set (DBCS) graphic data types.
- 10. On the Conversions tab, clear the Convert binary data (CCSID 65535) to text check box.
- 11. Click Apply, then OK. You are returned to the ODBC Data Source Administrator dialog.
- 12. Confirm that the new Data Source Name appears under User Data Sources.
- 13. Click **OK** to exit the ODBC configuration utility.
- **14.** From the Oracle GoldenGate directory on the target, run GGSCI and issue the DBLOGIN command to log into the target database. See *Reference for Oracle GoldenGate* for detailed syntax.

DBLOGIN SOURCEDB database, USERID db_user [, PASSWORD pw [encryption_options]]

Where:

- SOURCEDB database specifies the new data source name.
- USERID *db user*, PASSWORD *pw* are the Replicat database user profile and password.
- encryption options is optional password encryption.



Only BLOWFISH encryption is supported for DB2 for i systems.



Configuring Oracle GoldenGate for DB2 for i

This chapter contains instructions for configuring Oracle GoldenGate to capture source DB2 for i data and apply it to a supported target database.

Topics:

- · What to Expect from this Procedure
- Getting Started with Oracle GoldenGate
- Creating the Oracle GoldenGate Instance
- · Creating a GLOBALS File
- Creating a Data Definitions File
- Encrypting the Extract and Replicat Passwords
- Configuring Extract for Change Capture from DB2 for i
- Configuring Replicat for Change Delivery to DB2 for i
- · Next Steps in the Deployment
- When to Start Replicating Transactional Changes
- Testing Your Configuration

What to Expect from this Procedure

These instructions show you how to configure a set of basic Oracle GoldenGate parameter (configuration) files, one for each process that replicates transactional data changes from a DB2 for i source to a DB2 for i target, or to a different database type. Your business requirements probably will require a more complex topology, but this procedure forms a basis for the rest of your configuration steps.

This chapter focuses on the basic parameters that are specific to DB2 for i.

By performing these steps, you can:

- get the basic configuration files established.
- build upon them later by adding more parameters as you make decisions about features or requirements that apply to your environment.
- use copies of them to make additional parameter files faster than starting from scratch.

Getting Started with Oracle GoldenGate

Before proceeding with the configuration process, you should get familiar with the Oracle GoldenGate architecture, the command interface, and the methods for supplying input and instructions to the processes. See *Administering Oracle GoldenGate* for this information.

Creating the Oracle GoldenGate Instance

Each Oracle GoldenGate installation is rooted in the Manager process. This is the controller process that instantiates the Oracle GoldenGate processes, allocates port numbers, and performs file maintenance. Together, the Manager process and its child processes, and their related programs and files comprise an Oracle GoldenGate instance.

To run Oracle GoldenGate, a Manager process must be running on all systems that will be part of the Oracle GoldenGate environment. To run Manager, you first create a parameter file for it.

Creating a GLOBALS File

The GLOBALS parameter file contains parameters that affect all processes within an Oracle GoldenGate instance. The GLOBALS parameter NAMECCSID is specific to DB2 for i and may be required if the SQL catalog contains object names that are referenced by a different CCSID than the system CCSID. The SQL catalog is created in the system CCSID and does not indicate this difference when queried. Oracle GoldenGate makes queries to the catalog and could retrieve the name incorrectly unless NAMECCSID is used to supply the correct CCSID value. For more information, see *Reference for Oracle GoldenGate*.

Creating a Data Definitions File

When replicating data from one table to another, an important consideration is whether the column structures (metadata) of the source and target tables are identical. Oracle GoldenGate looks up metadata for the following purposes:

- On the source, to supply complete information about captured operations to the Replicat process.
- On the target, to determine the structures of the target tables, so that the replicated data is correctly mapped and converted (if needed) by Replicat.

When source and target table definitions are dissimilar, Oracle GoldenGate must perform a conversion from one format to the other. To perform conversions, both sets of definitions must be known to Oracle GoldenGate. Oracle GoldenGate can query the local database to get one set of definitions, but it must rely on a *data-definitions file* to get definitions from the remote database. The data-definitions file contains information about the metadata of the data that is being replicated.

To create a definitions file, you configure and run the DEFGEN utility and then transfer the definitions file to the target system. This file must be in place on the target system before you start the Oracle GoldenGate processes for the first time.

Encrypting the Extract and Replicat Passwords

It is strongly recommended that you encrypt the passwords of the user profiles that will be used for the primary and data pump Extracts, and for the Replicat process. Oracle GoldenGate must use Blowfish encryption on the DB2 for i platform. The standard Oracle GoldenGate encryption method of AES (Advanced Encryption Standard) is supported by the IBM i platform. To encrypt the password, see Working with Runtime Parameters in *Administering Oracle GoldenGate*. It also contains information about how to encrypt data within disk storage and across TCP/IP.





The Oracle GoldenGate credential store is not supported by the iSeries platform.

Configuring Extract for Change Capture from DB2 for i

Perform these steps on the source system to configure the primary Extract and the data pump Extract that support change capture and transport across the network.

- Configuring the Primary Extract
- Configuring the Data Pump

Configuring the Primary Extract

These steps configure the primary Extract to capture transaction data from a source DB2 LUW and write the data to a local trail for temporary storage.

1. In GGSCI on the source system, create the Extract parameter file.

```
EDIT PARAMS name
```

Where: name is the name of the primary Extract.

Enter the Extract parameters in the order shown, starting a new line for each parameter statement.

Basic parameters for the primary Extract

EXTRACT finance
SOURCEDB mysource, USERIDALIAS myalias
ENCRYPTTRAIL AES192
EXTTRAIL /ggs/dirdat/lt
TABLE hr.*;

Parameter	Description	
EXTRACT group	group is the name of the Extract group.	
SOURCEDB database, USERIDALIAS alias	Specifies the real name of the source DB2 for i database (not an alias), plus the alias of the database login credential of the user that is assigned to Extract. This credential must exist in the Oracle GoldenGate credential store. For more information, see Database User for Oracle GoldenGate Processes.	
ENCRYPTTRAIL algorithm	Encrypts the local trail.	
EXTTRAIL pathname	Specifies the path name of the local trail to which the primary Extract writes captured data for temporary storage.	



Parameter	Description	
TABLE schema.object;	Specifies the database object for which to capture data.	
	TABLE is a required keyword.	
	 schema is the schema name or a wildcarded set of schemas. 	
	 object is the table name, or a wildcarded set of tables. 	
	The question mark (?) wildcard is not supported for this database. Note that only the asterisk (*) wildcard is supported for DB2 LUW.	
	Terminate the parameter statement with a semi-colon.	
	To exclude tables from a wildcard specification, use the TABLEEXCLUDE parameter.	
	For more information and for additional options that control data filtering, mapping, and manipulation, see TABLE MAP in <i>Reference for Oracle GoldenGate</i> .	

- 3. Enter any optional Extract parameters that are recommended for your configuration. You can edit this file at any point before starting processing by using the EDIT PARAMS command in GGSCI.
- 4. Save and close the file.

Configuring the Data Pump

These steps configure the data pump that reads the local trail and sends the data across the network to a remote trail.

1. In GGSCI on the source system, create the data-pump parameter file.

EDIT PARAMS name

Where: name is the name of the data pump Extract group.

Enter the data-pump parameters in the order shown, starting a new line for each parameter statement. Your input variables will be different. See the following table for description.

Basic parameters for the data-pump Extract group:

```
EXTRACT extpump
SOURCEDB FINANCE USERID ogg, PASSWORD AACAAAAAAAA, BLOWFISH ENCRYPTKEY mykey
RMTHOST fin1, MGRPORT 7809
RMTTRAIL /ggs/dirdat/rt
TABLE hr.*;
```

Parameter	Description	
EXTRACT group	group name is the name of the data pump.	
SOURCEDB database USERID user, PASSWORD password, BLOWFISH ENCRYPTKEY keyname	 Specifies database connection information. SOURCEDB specifies the <i>default</i> DB 2 for i database. USERID specifies the Extract database user profile. PASSWORD specifies the user's password that was encrypted with the ENCRYPT PASSWORD command. Enter or paste the encrypted password after the PASSWORD keyword. BLOWFISH ENCRYPTKEY <i>keyname</i> specifies the name of the lookup key in the local ENCKEYS file. 	
DECRYPTTRAIL BLOWFISH	Decrypts the input trail.	



Parameter	Description	
RMTHOST hostname, MGRPORT portnumber	 RMTHOST specifies the name or IP address of the target system. MGRPORT specifies the port number where Manager is running on the target. 	
ENCRYPTTRAIL BLOWFISH	Encrypts the remote trail with Blowfish encryption.	
RMTTRAIL pathname	Specifies the path name of the remote trail.	
TABLE owner.table;	Specifies a table or tables to process. Terminate the TABLE statement with a semi-colon. To exclude tables from a wildcard specification, use the TABLEEXCLUDE owner.table	

- 3. Enter any optional Extract parameters that are recommended elsewhere in this manual and any others shown in Summary of Extract Parameters in *Reference for Oracle GoldenGate*.
- Save and close the file.

Configuring Replicat for Change Delivery to DB2 for i

These steps configure Replicat to apply data to a DB2 for i target database, operating either on the target system or on a remote Windows or Linux system. To configure Replicat for change delivery to a different database type, such as an Oracle database, follow the directions in the Oracle GoldenGate Installation and Configuration guide for that database. There may be additional parameters and requirements for delivery to that database type.



There does not have to be a database on a Windows or Linux machine to support connection by ODBC by Replicat.

- Creating a Checkpoint Table
- · Configuring Replicat

Creating a Checkpoint Table

Replicat maintains its checkpoints in a checkpoint table in the DB2 for i target database. Each checkpoint is written to the checkpoint table, that must be journaled, within the Replicat transaction. Because a checkpoint either succeeds or fails with the transaction, Replicat ensures that a transaction is only applied once, even if there is a failure of the process or the database.

A common method of create the checkpoint table with journaling is as follows:

1. In GGSCI on the target system, create the Replicat checkpoint file.

```
DEFAULTJOURNAL library_name/journal_name
```

Where: <code>library_name</code> is the name of the library and <code>journal_name</code> is the name of the default journal.

2. Add the checkpoint table.

ADD CHECKPOINTTABLE $library_name.$ chkptab

Successfully created checkpoint table kgr.chkptab

Add journaling to the checkpoint table.

ADD TRANDATA library name. CHKPTAB

For more information about creating a checkpoint table, see Administering Oracle GoldenGate.

Configuring Replicat

These steps configure the Replicat process in a basic way without any special mapping or conversion of the data.

1. In GGSCI on the target system, create the Replicat parameter file.

EDIT PARAMS name

Where: name is the name of the Replicat group.

Enter the Replicat parameters in the order shown, starting a new line for each parameter statement.

REPLICAT financer

TARGETDB FINANCIAL USERID ogg, PASSWORD AACAAAAAAAA, BLOWFISH ENCRYPTKEY mykey ASSUMETARGETDEFS

- -- Instead of ASSUMETARGETDEFS, use SOURCEDEFS if replicating from
- -- DB2 LUW to a different database type, or from a DB2 DB2 LUW source
- -- that is not identical in definitions to a target DB2 LUW database.
- -- SOURCEDEFS /users/ogg/dirdef/defsfile

DISCARDFILE /users/ogg/disc

MAP hr.*, TARGET hr2.*;

Parameter	Description	
REPLICAT group	group is the name of the Replicat group.	
TARGETDB database USERID user, PASSWORD password, BLOWFISH ENCRYPTKEY keyname	 Specifies database connection information. SOURCEDB specifies the data source name (DSN) of the target DB2 LUW database. USERID specifies the Replicat database user profile. PASSWORD specifies the user's password that was encrypted with the ENCRYPT PASSWORD command. Enter or paste the encrypted password after the PASSWORD keyword. BLOWFISH ENCRYPTKEY keyname specifies the name of the lookup key in the local ENCKEYS file. 	
DECRYPTTRAIL BLOWFISH	Decrypts the input trail.	
SOURCEDEFS pathname ASSUMETARGETDEFS	Specifies how to interpret data definitions. Use <code>SOURCEDEFS</code> if the source and target tables have different definitions, such as when replicating data between dissimilar IBM databases or from an IBM database to an Oracle database. For <code>pathname</code> , specify the source datadefinitions file that you created with the <code>DEFGEN</code> utility. Use <code>ASSUMETARGETDEFS</code> if the source and target tables are all <code>DB2</code> LUW and have the same definitions.	



Parameter	Description
MAP owner.table, TARGET owner.table;	Specifies a relationship between a source and target table or tables. The MAP clause specifies the source objects, and the TARGET clause specifies the target objects to which the source objects are mapped.
	• owner is the schema or library name.
	• table is the name of a table or a wildcard definition for multiple tables.
	Terminate the MAP statement with a semi-colon.
	To exclude tables from a wildcard specification, use the MAPEXCLUDE parameter.
	For more information and for additional options that control data filtering, mapping, and manipulation, see MAP in <i>Reference for Oracle GoldenGate</i> .

- 3. Enter any optional Extract parameters that are recommended elsewhere in this manual and any others shown in Summary of Extract Parameters.
- 4. Save and close the file.

Next Steps in the Deployment

Because of its flexibility, Oracle GoldenGate offers numerous features and options that must be considered before you start any processes. To further configure Oracle GoldenGate to suit your business needs, see the following:

- For additional configuration guidelines to achieve specific replication topologies, see Administering Oracle GoldenGate. This guide also contains information about:
 - Oracle GoldenGate architecture
 - Oracle GoldenGate commands
 - Oracle GoldenGate initial load methods
 - Using customization features
 - Mapping columns that contain dissimilar data
 - Data filtering and manipulation
- For syntax options and descriptions of Oracle GoldenGate GGSCI commands and Oracle GoldenGate parameters shown in this guide, see *Reference for Oracle GoldenGate*.

When to Start Replicating Transactional Changes

You must start replication when the source and target data is in a synchronized state, where the corresponding rows in the source and target tables contain identical data values. Unless you are starting with brand new source and target databases with no current user activity, you will need to activate change capture and apply processes to handle ongoing transactional changes while an initial load is being applied to the target. This process is known as *initial synchronization*, or also as *instantiation*. The initial load captures a point-in-time snapshot of the source data and applies it to the target, while Oracle GoldenGate maintains any changes that are made after that point.

See Instantiating Oracle GoldenGate with an Initial Load in *Administering Oracle GoldenGate* for instantiation options.

- Starting Extract During Instantiation
- Changing the Position of Extract to a Later Time



Starting Extract During Instantiation

When Extract starts for the first time to begin capturing data during the instantiation process, it captures all of the transaction data that it encounters after the specified start point, but none of the data that occurred before that point. To ensure that Extract does not start in the middle of ongoing transactions that would be discarded, set the tables that are to be captured to an inactive state. You can either put the system into a restricted state by using the ALCOBJ command to lock the objects or libraries, or you can force all of the current transactions on those tables to stop at a certain point.

After initialization is complete, remember to unlock any objects that you locked. To do so, log off of the session that locked the objects or use the DLCOBJ command from the OS/400 command line.

Changing the Position of Extract to a Later Time

You may at some point, over the life of an Extract run, need to set the position of Extract in the data stream manually. To reposition Extract, use the ALTER EXTRACT command in GGSCI. To help you identify any given Extract read position, the INFO EXTRACT command shows the positions for each journal in an Extract configuration, including the journal receiver information. See *Reference for Oracle GoldenGate* to know more.



Because the journals can have a transaction split among them, if a given journal is independently repositioned far into the past, the resulting latency from reprocessing the entries may cause the already-read journals to stall until the reading of the latent journal catches up.

Testing Your Configuration

It is important to test your configuration in a test environment before deploying it live on your production machines. This is especially important in an active-active or high availability configuration, where trusted source data may be touched by the replication processes. Testing enables you to find and resolve any configuration mistakes or data issues without the need to interrupt user activity for re-loads on the target or other troubleshooting activities.



7

Instantiating and Starting Oracle GoldenGate Replication

This chapter contains instructions for configuring an initial load of target data, adding the required processes to instantiate replication, and perform the instantiation. The expected outcome of these steps is that source-target data is made consistent (known as the initial synchronization), and that Oracle GoldenGate captures and delivers ongoing transactional changes so that consistency is maintained going forward.

Topics:

- About the Instantiation Process
- Overview of Basic Oracle GoldenGate Instantiation Steps
- · Satisfying Prerequisites for Instantiation
- Making the Instantiation Procedure More Efficient
- Configuring the Initial Load
- Adding Change-Capture and Change-Delivery processes
- Performing the Target Instantiation
- Monitoring Processing after the Instantiation
- Backing up Your Oracle GoldenGate Environment
- Positioning Extract After Startup

About the Instantiation Process

During the initialization of the Oracle GoldenGate environment, you will be doing an initial data synchronization and starting the Oracle GoldenGate processes for the first time. In conjunction with those procedures, you will be creating the process groups for which you created parameter files in Configuring Oracle GoldenGate for DB2 for i.

To create an Extract process group, an initial start position for data capture must be established. This initial position will be based on a transaction boundary that is based on either of the following:

- a timestamp
- the end of the journal(s)
- A specific system sequence number
- A specific sequence number in the journal(s)

When Extract starts for the first time to begin capturing data, it captures all of the transaction data that it encounters after the specified start point, but none of the data that occurred before that point. To ensure that Extract does not start in the middle of ongoing transactions that would be discarded, set the tables that are to be captured to an inactive state. You can either put the system into a restricted state by using the ALCOBJ command to lock the objects or libraries, or you can force all of the current transactions on those tables to stop at a certain point.

After initialization is complete, remember to unlock any objects that you locked. To do so, log off of the session that locked the objects or use the DLCOBJ command from the OS/400 command line.

Overview of Basic Oracle GoldenGate Instantiation Steps

These instructions show you how to instantiate the basic replication environment that you configured in Chapter 4. These steps are:

- Satisfying Prerequisites for Instantiation
- Making the Instantiation Procedure More Efficient
- Configuring the Initial Load
- Adding Change-Capture and Change-Delivery processes
- Performing the Target Instantiation
- Monitoring Processing after the Instantiation
- Backing up Your Oracle GoldenGate Environment
- Positioning Extract After Startup

Satisfying Prerequisites for Instantiation

These steps must be taken before starting any Oracle GoldenGate processes or native database load processes.

- Configure Change Capture and Delivery
- Add Collision Handling
- Prepare the Target Tables

Configure Change Capture and Delivery

By the time you are ready to instantiate the replication environment, all of your Extract and Replicat process groups must be configured with completed parameter files as directed in "Configuring Oracle GoldenGate for DB2 for i".

In addition, all of the other setup requirements in this manual must be satisfied.

Add Collision Handling

If the source database will remain active during the initial load, collision-handling logic must be added to the Replicat parameter file. This logic handles conflicts that occur because static data is being loaded to the target tables while Oracle GoldenGate replicates transactional changes to those tables.

To handle collisions, add the HANDLECOLLISIONS parameter to the Replicat parameter file to resolve:

- INSERT operations for which the row already exists.
- UPDATE and DELETE operations for which the row does not exist.

For more information about this parameter, see the Oracle GoldenGate *Windows and UNIX Reference Guide*.



Prepare the Target Tables

The following are suggestions that can make the load go faster and help you to avoid errors.

- Data: Make certain that the target tables are empty. Otherwise, there may be duplicate-row errors or conflicts between existing rows and rows that are being loaded.
- Constraints: If you have not done so already, disable foreign-key constraints and check constraints. Foreign-key constraints can cause errors, and check constraints can slow down the loading process.
- Indexes: Remove indexes from the target tables. Indexes are not necessary for the inserts
 performed by the initial load process and will slow it down significantly. You can add back
 the indexes after the load is finished.
- Keys: To use the HANDLECOLLISIONS function to reconcile incremental data changes with the load, each target table must have a primary or unique key. If you cannot create a key through your application, use the KEYCOLS option of the TABLE and MAP parameters to specify columns as a substitute key for Oracle GoldenGate's purposes. If you cannot create keys, the affected source table must be quiesced for the load.

Making the Instantiation Procedure More Efficient

The following are some suggestions for making the instantiation process move more efficiently.

- Share Parameters Between Process Groups
- Use Parallel Processes

Share Parameters Between Process Groups

Some of the parameters that you use in a change-synchronization parameter file also are required in an initial-load Extract and initial-load Replicat parameter file. To take advantage of the commonalities, you can use any of the following methods:

- Copy common parameters from one parameter file to another.
- Store the common parameters in a central file and use the OBEY parameter in each parameter file to retrieve them.
- Create an Oracle GoldenGate macro for the common parameters and then call the macro from each parameter file with the MACRO parameter.

Use Parallel Processes

You can configure parallel initial-load processes to perform the initial load more quickly. It is important to keep tables with foreign-key relationships within the same set of processes. You can isolate large tables from smaller ones by using different sets of processes, or simply apportion the load across any number of process sets. To configure parallel processes correctly, see Administering Oracle GoldenGate for Windows and UNIX.

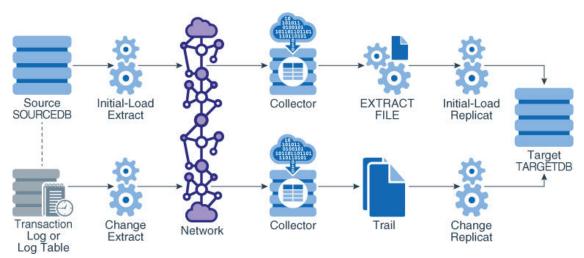
Configuring the Initial Load

Oracle GoldenGate supports the following load methods specifically for Oracle:

Configuring an Initial Load from File to Replicat

Configuring an initial load with a database utility

Configuring an Initial Load from File to Replicat



To use Replicat to establish the target data, you use an initial-load Extract to extract source records from the source tables and write them to an extract file in canonical format. From the file, an initial-load Replicat loads the data using the database interface. During the load, the change-synchronization groups extract and replicate incremental changes, which are then reconciled with the results of the load.

During the load, the records are applied to the target database one record at a time, so this method may be considerably slower than using a native DB2 for i load utility. This method permits data transformation to be done on either the source or target system.

To Configure a Load from File to Replicat

1. On the source and target systems, run GGSCI and start Manager.

START MANAGER

2. On the source system, issue the following command to create an initial-load Extract parameter file. This Extract should have a different name from the Extract groups that capture the transactional data.

EDIT PARAMS initial-load Extract name

3. Enter the parameters listed in the following table in the order shown, starting a new line for each parameter statement.

Parameter	Description
SOURCEISTABLE	Designates Extract as an initial load process extracting records directly from the source tables.



Parameter	Description
SOURCEDB database USERID user id, PASSWORD password, BLOWFISH ENCRYPTKEY keyname	Specifies database connection information. SOURCEDB specifies the name of the source database
	 USERID specifies the Extract database user profile. PASSWORD specifies the user's password that was encrypted with the ENCRYPT PASSWORD command (see "Encrypting the Extract and Replicat Passwords"). Enter or paste the encrypted password after the PASSWORD keyword. BLOWFISH ENCRYPTKEY keyname specifies the name of the lookup key in the local ENCKEYS file.
RMTHOST hostname, MGRPORT portnumber, [encryption options]	 RMTHOST specifies the name or IP address of the target system. MGRPORT specifies the port number where Manager is running on the target. encryption options specifies optional encryption of data across TCP/IP. For additional options and encryption details, see Reference for Oracle GoldenGate for Windows and UNIX.
ENCRYPTTRAIL BLOWFISH KEYNAME keyname	Encrypts the remote file with Blowfish encryption. For more information about security, see Administering Oracle GoldenGate for Windows and UNIX.
RMTFILE path name, [MEGABYTES n]	Specifies the remote file to which the load data will be written. Oracle GoldenGate creates this file during the load.
 path name is the relative or fully qualified name of the file. MEGABYTES designates the size of each file. 	Note: The size of an extract file cannot exceed 2GB.
TABLE owner.table;	Specifies a source table or tables for initial data extraction.
 owner is the library or schema name. table is the name of the table or a group of tables defined with wildcards. To exclude tables from a wildcard specification, use the TABLEEXCLUDE parameter. 	

- - 4. Enter any appropriate optional Extract parameters listed in Reference for Oracle GoldenGate for Windows and UNIX.
 - 5. Save and close the parameter file.
 - 6. On the target system, issue the following command to create an initial-load Replicat parameter file. This Replicat should have a different name from the Replicat group that applies the transactional data.

EDIT PARAMS initial-load Replicat name

7. Enter the parameters listed in Table 7-1 in the order shown, starting a new line for each parameter statement.

Table 7-1 Initial Load Replicat Parameters for Loading Data from File to Replicat

Parameter	Description
SPECIALRUN	Implements the initial-load Replicat as a one-time run that does not use checkpoints.

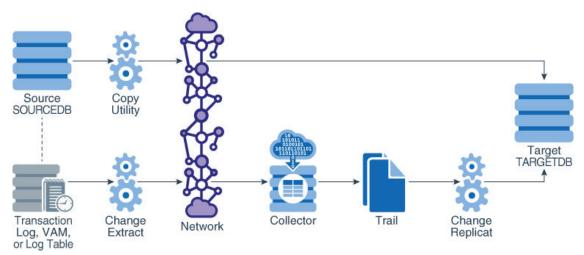


Table 7-1 (Cont.) Initial Load Replicat Parameters for Loading Data from File to Replicat

Parameter	Description
END RUNTIME	Directs the initial-load Replicat to terminate when the load is finished.
TARGETDB database, USERID user id, PASSWORD pw, algorithm ENCRYPTKEY keyname	 Specifies database connection information. TARGETDB specifies the Data Source Name that is defined for the DB2 for i target database through the ODBC interface on the Windows or Linux system. USERID specifies the Replicat database user profile. PASSWORD specifies the user's password that was encrypted with the ENCRYPT PASSWORD command (see "Encrypting the Extract and Replicat Passwords"). Enter or paste the encrypted password after the PASSWORD keyword.
	 algorithm ENCRYPTKEY keyname specifies the encryption method and keyname that was specified in the ENCRYPT PASSWORD command.
DECRYPTTRAIL BLOWFISH KEYNAME keyname	Decrypts the input trail. BLOWFISH is required because this is the algorithm that is supported to encrypt the file from DB2 for i.
EXTFILE path name EXTTRAIL path name	Specifies the input extract file specified with the Extract parameter RMTFILE.
 path name is the relative or fully qualified name of the file or trail. Use EXTTRAIL only if you used the MAXFILES option of the RMTFILE parameter in the Extract parameter file. 	
SOURCEDEFS file name ASSUMETARGETDEFS	Specifies how to interpret data definitions.
 Use SOURCEDEFS if the source and target tables have different definitions. Specify the relative or fully qualified name of the source-definitions file generated by the DEFGEN utility. 	
 Use ASSUMETARGETDEFS if the source and target tables have the same definitions. 	
MAP owner.table, TARGET owner.table;	Specifies a relationship between a source and target table or tables.
 owner is the schema name. table is the name of a table or a wildcard definition for multiple tables. To exclude tables from a wildcard specification, use the MAPEXCLUDE parameter. 	

- 8. Enter any appropriate optional Replicat parameters listed in the Reference for Oracle GoldenGate for Windows and UNIX.
- 9. Save and close the file.

Configuring an initial load with a database utility



This graphic shows the parallel flows of the initial load and the ongoing capture and replication of transactional changes during the load period. The copy utility writes the data to a file, which is loaded to the target. Meanwhile, an Extract process captures change data and sends it to a trail on the target for Replicat to read and apply to the target.

For an initial load between two DB2 for i source and target systems, you can use the DB2 for i system utilities to establish the target data. To do this, you save the file(s) that you want to load to the target by using the SAVOBJ or SAVLIB commands, and then you restore them on the target using the RSTOBJ or RSTLIB commands.

Another alternative is to use the DB2 for i commands CPYTOIMPF (Copy to Import File) and CPYFRMIMPF (Copy from Import File) to create files that can be used with the bulk load utilities of other databases. See the DB2 for i Information Center documentation for more details on "Copying between different systems."

In both cases, no special configuration of any Oracle GoldenGate initial-load processes is needed. You use the change-synchronization process groups that you configured in Configuring Oracle GoldenGate for DB2 for i . You start a change-synchronization Extract group to extract ongoing data changes while you are making the copy and loading it. When the copy is finished, you start the change-synchronization Replicat group to re-synchronize rows that were changed while the copy was being applied. From that point forward, both Extract and Replicat continue running to maintain data synchronization. See "Adding Change-Capture and Change-Delivery processes".

Adding Change-Capture and Change-Delivery processes



Perform these steps at or close to the time that you are ready to start the initial load and change capture.

These steps establish the Oracle GoldenGate Extract, data pump, and Replicat processes that you configured in Configuring Oracle GoldenGate for DB2 for i . Collectively known as the "change-synchronization" processes, these are the processes that:

- capture and apply ongoing source changes while the load is being performed on the target
- reconcile any collisions that occur

Note:

Perform these steps as close as possible to the time that you plan to start the initial load processes. You will start these processes during the initial load steps.

- Add the Primary Extract
- Add the Local Trail
- Add the Data Pump Extract Group
- Add the Remote Trail
- Add the Replicat Group

Add the Primary Extract

These steps add the primary Extract that captures change data.

- Understanding the Primary Extract Start Point
- Establishing the Required and Optional Extract Start Points

Understanding the Primary Extract Start Point

When you add the primary Extract group, you establish an initial start position for data capture. This initial position can be a transaction boundary that is based on either of the following:

- a timestamp
- the end of the journal(s)
- a specific system sequence number
- a specific journal sequence number (per journal)

The options that are available to you assume a global start point and optional journal-specific start points.

- To position by a timestamp, at the end of the journals, or at a system sequence number, you will use the ADD EXTRACT command with the appropriate option. This command establishes a global start point for all journals and is a required first step.
- After issuing the ADD EXTRACT command, you can then optionally position any *specific* journal at a specific journal sequence number by using the ALTER EXTRACT command with an appropriate journal option.

Establishing the Required and Optional Extract Start Points

These steps include the ADD EXTRACT and ALTER EXTRACT commands to enable you to establish your desired start points.

- 1. Run GGSCI.
- 2. Issue the ADD EXTRACT command to add the primary Extract group and establish the global start point.



```
ADD EXTRACT group name, TRANLOG {

, BEGIN {NOW | yyyy-mm-dd[hh:mi:[ss[.cccccc]]]} |
, EOF |
, SEQNO seqno
}
```

Where:

- group name is the name of the primary Extract group that captures the transactional changes.
- TRANLOG specifies the journals as the data source.
- BEGIN specifies to begin capturing data as of a specific *time*. Select one of two options: NOW starts at the first record that is timestamped at the same time that BEGIN is issued. yyyy-mm-dd[hh:mi:[ss[.ccccc]]] starts at an explicit timestamp. Logs from this timestamp must be available.
- SEQNO segno specifies to begin capturing data at, or just after, a system sequence number, which is a decimal number up to 20 digits in length.
- 3. (Optional) Issue the following command to alter any ADD EXTRACT start position to set the start position for a specific journal in the same Extract configuration. A specific journal position set with ALTER EXTRACT does not affect any global position that was previously set with ADD EXTRACT or ALTER EXTRACT; however a global position set with ALTER EXTRACT overrides any specific journal positions that were previously set in the same Extract configuration.

```
ALTER EXTRACT group name,
{
ALTER EXTRACT {BEGIN {NOW | yyyy-mm-dd [hh:mi:[ss[.ccccc]]] [JOURNAL journal_library/journal_name [[JRNRCV receiver_library/receiver_name]] |
, EOF [JOURNAL journal_library/journal_name [[JRNRCV receiver_library/receiver_name]] |
, SEQNO seqno [JOURNAL journal_library/journal_name [[JRNRCV receiver_library/receiver_name]]
}
```



SEQNO, when used with a journal in ALTER EXTRACT, is the journal sequence number that is relative to that specific journal, not the system sequence number that is global across journals.

Example 7-1 Timestamp Start Point

ADD EXTRACT finance, TRANLOG, BEGIN 2011-01-01 12:00:00.000000

Example 7-2 NOW Start Point

ADD EXTRACT finance, TRANLOG, BEGIN NOW

Example 7-3 System Sequence Number Start Point

ADD EXTRACT finance, TRANLOG, SEQNO 2952



Example 7-4 Journal Start Point

ALTER EXTRACT finance, SEQNO 1234 JOURNAL accts/acctsjrn

Example 7-5 Journal and Receiver Start Point

ALTER EXTRACT finance, SEQNO 1234 JOURNAL accts/acctsjrn JRNRCV accts/jrnrcv0005

Add the Local Trail

This step adds the local trail to which the primary Extract writes captured data.

In GGSCI on the source system, issue the ADD EXTTRAIL command:

ADD EXTTRAIL pathname, EXTRACT group name

Where:

- EXTTRAIL specifies that the trail is to be created on the local system.
- pathname is the relative or fully qualified name of the trail, including the two-character name
- EXTRACT group name is the name of the primary Extract group.

Example 7-6

ADD EXTTRAIL /ggs/dirdat/lt, EXTRACT finance

Add the Data Pump Extract Group

This step adds the data pump that reads the local trail and sends the data to the target.

In GGSCI on the source system, issue the ADD EXTRACT command.

ADD EXTRACT group name, EXTTRAILSOURCE trail name

Where:

- group name is the name of the data-pump Extract group.
- EXTTRAILSOURCE trail name is the relative or fully qualified name of the local trail.

Example 7-7

ADD EXTRACT financep, EXTTRAILSOURCE c:\ggs\dirdat\lt

Add the Remote Trail

This step adds the remote trail. Although it is read by Replicat, this trail must be associated with the data pump, so it must be added on the source system, not the target.

In GGSCI on the source system, issue the following command:

ADD RMTTRAIL pathname, EXTRACT group name

Where:

- RMTTRAIL specifies that the trail is to be created on the target system, and pathname is the relative or fully qualified name of the trail, including the two-character name.
- EXTRACT group name is the name of the data-pump Extract group.



Example 7-8

ADD RMTTRAIL /ggs/dirdat/rt, EXTRACT financep

Add the Replicat Group

These steps add the Replicat group that reads the remote trail (which gets created automatically on the target) and applies the data changes to the target Oracle database.

- Run GGSCI on the target system.
- 2. Issue the ADD REPLICAT command.

ADD REPLICAT group name, EXTTRAIL pathname

Where:

- group name is the name of the Replicat group.
- EXTTRAIL pathname is the relative or fully qualified name of the remote trail, including the two-character name.

Example 7-9

ADD REPLICAT financer, EXTTRAIL c:\ggs\dirdat\rt

Performing the Target Instantiation

This procedure instantiates the target tables while Oracle GoldenGate captures ongoing transactional changes on the source and stores them until they can be applied on the target. By the time you perform the instantiation of the target tables, the entire Oracle GoldenGate environment should be configured for change capture and delivery, as should the initial-load processes if using Oracle GoldenGate as an initial-load utility.

- To Perform Instantiation from File to Replicat
- To Perform Instantiation with a Database Utility

To Perform Instantiation from File to Replicat

- 1. Make certain that you have addressed the requirements in Satisfying Prerequisites for Instantiation.
- 2. On the source and target systems, run GGSCI and start the Manager process.

```
START MANAGER
```

On the source system, start the primary and data pump Extract groups to start change extraction.

```
START EXTRACT primary Extract group name
START EXTRACT data pump Extract group name
```

4. From the directory where Oracle GoldenGate is installed on the source system, start the initial-load Extract as follows:

```
\ /GGS directory/extract paramfile dirprm/initial-load Extract name.prm reportfile path name
```

Where: initial-load Extract name is the name of the initial-load Extract that you used when creating the parameter file, and path name is the relative or fully qualified name of

- the Extract report file (by default the dirrpt sub-directory of the Oracle GoldenGate installation directory).
- Verify the progress and results of the initial extraction by viewing the Extract report file using the operating system's standard method for viewing files.
- 6. Wait until the initial extraction is finished.
- 7. On the target system, start the initial-load Replicat.

```
$ /GGS directory/replicat paramfile dirprm/initial-load Replicat name.prm
reportfile path name
```

Where: initial-load Replicat name is the name of the initial-load Replicat that you used when creating the parameter file, and path name is the relative or fully qualified name of the Replicat report file (by default the dirrpt sub-directory of the Oracle GoldenGate installation directory).

- 8. When the initial-load Replicat is finished running, verify the results by viewing the Replicat report file using the operating system's standard method for viewing files.
- 9. On the target system, start change replication.

```
START REPLICAT Replicat group name
```

10. On the target system, issue the following command to verify the status of change replication.

```
INFO REPLICAT Replicat group name
```

- 11. Continue to issue the INFO REPLICAT command until you have verified that Replicat posted all of the change data that was generated during the initial load. For example, if the initialload Extract stopped at 12:05, make sure Replicat posted data up to that point.
- 12. On the target system, issue the following command to turn off the HANDLECOLLISIONS parameter and disable the initial-load error handling.

```
SEND REPLICAT Replicat group name, NOHANDLECOLLISIONS
```

13. On the target system, edit the Replicat parameter file to remove the HANDLECOLLISIONS parameter. This prevents HANDLECOLLISIONS from being enabled again the next time Replicat starts.



Caution:

Do not use the VIEW PARAMS or EDIT PARAMS command to view or edit an existing parameter file that is in a character set other than that of the local operating system (such as one where the CHARSET option was used to specify a different character set). View the parameter file from outside GGSCI if this is the case; otherwise, the contents may become corrupted.

14. Save and close the parameter file.

From this point forward, Oracle GoldenGate continues to synchronize data changes.

To Perform Instantiation with a Database Utility

- Make certain that you have addressed the requirements in "Satisfying Prerequisites for Instantiation".
- On the source and target systems, run GGSCI and start the Manager process.

START MANAGER

On the source system, start the primary and data pump Extract groups to start change extraction.

```
START EXTRACT primary Extract group name
START EXTRACT data pump Extract group name
```

- 4. On the source system, start making the copy.
- 5. Wait until the copy is finished and record the time of completion.
- 6. View the Replicat parameter file to make certain that the HANDLECOLLISIONS parameter is listed. If not, edit the file and add the parameter to the file.

EDIT PARAMS Replicat group name



Do not use the VIEW PARAMS or EDIT PARAMS command to view or edit an existing parameter file that is in a character set other than that of the local operating system (such as one where the CHARSET option was used to specify a different character set). View the parameter file from outside GGSCI if this is the case; otherwise, the contents may become corrupted.

7. On the target system, start change replication.

START REPLICAT Replicat group name

8. On the target system, issue the following command to verify the status of change replication.

INFO REPLICAT Replicat group name

- 9. Continue to issue the INFO REPLICAT command until you have verified that change replication has posted all of the change data that was generated during the initial load. Reference the time of completion that you recorded. For example, if the copy stopped at 12:05, make sure change replication has posted data up to that point.
- **10.** On the target system, issue the following command to turn off the HANDLECOLLISIONS parameter and disable the initial-load error handling.

SEND REPLICAT Replicat group name, NOHANDLECOLLISIONS

11. On the target system, edit the Replicat parameter file to remove the HANDLECOLLISIONS parameter. This prevents HANDLECOLLISIONS from being enabled again the next time Replicat starts.



Caution:

Do not use the VIEW PARAMS or EDIT PARAMS command to view or edit an existing parameter file that is in a character set other than that of the local operating system (such as one where the CHARSET option was used to specify a different character set). View the parameter file from outside GGSCI if this is the case; otherwise, the contents may become corrupted.

12. Save and close the parameter file.

From this point forward, Oracle GoldenGate continues to synchronize data changes.



Monitoring Processing after the Instantiation

After the target is instantiated and replication is in effect, you should view the status, lag, and overall health of the replication environment to ensure that processes are running properly, that there are no warnings in the Oracle GoldenGate error log, and that lag is at an acceptable level. You can view Oracle GoldenGate processes from:

- GGSCI: For information about monitoring processes, see Administering Oracle GoldenGate for Windows and UNIX.
- Oracle GoldenGate Monitor: See the administration documentation and online help for that product. Oracle GoldenGate Monitor provides a graphical-based monitoring environment for all of your Oracle GoldenGate instances.

You also should verify that capture and delivery is being performed for all of the tables in the Oracle GoldenGate configuration, and that the source and target data are synchronized. You can use the Oracle GoldenGate Veridata product for this purpose.

Backing up Your Oracle GoldenGate Environment

After you start Oracle GoldenGate processing, an effective backup routine is critical to preserving the state of processing in the event of a failure. Unless the Oracle GoldenGate working files can be restored, the entire replication environment must be re-instantiated, complete with new initial loads.

As a best practice, include the entire Oracle GoldenGate home installation in your backup routines. This directory contains critical sub-directories, files and programs. The most critical working files in this directory consume the vast majority of backup space; therefore it makes sense just to back up the entire installation directory for fast, simple recovery.

Positioning Extract After Startup

You may at some point, over the life of an Extract run, need to set the position of Extract in the data stream manually. To reposition Extract, use the ALTER EXTRACT command in GGSCI. To help you identify any given Extract read position, the INFO EXTRACT command shows the positions for each journal in an Extract configuration, including the journal receiver information. For more information about these commands, see Reference for Oracle GoldenGate for Windows and UNIX.



Because the extract will be synchronizing all of the journals in the extract by system sequence number because it is possible for a transaction to be split across them, if a given journal is independently repositioned far into the past, the resulting latency from reprocessing the entries will cause the already-read journals to stall until the reading of the latent journal catches up.



Using Remote Journal

This chapter contains instructions for remote journal preparation and adding a remote journal. Remote Journal support in the IBM DB2 for i operating system provides the ability for a system to replicate, in its entirety, a sequence of journal entries from one DB2 for i system to another. Once setup, this replication is handled automatically and transparently by the operating system. The entries that are replicated are placed in a journal on the target system that is available to be read by an application in the same way as on the source system. You must have an understanding of how to setup and use remote journaling on an DB2 for i system to use this feature with Oracle GoldenGate. There are no special software requirements for either Oracle GoldenGate or the DB2 for i systems to use remote journaling.

Topics:

- Preparing to Use Remote Journals
- Adding a Remote Journal

Preparing to Use Remote Journals

Before establishing the remote journal environment, complete the following steps:

- 1. Determine the extent of your remote journal network or environment.
- Library redirection is the ability to allow the remote journal and associated journal receivers
 to reside in different libraries on the target system from the corresponding source journal
 and its associated journal receivers.
 - Determine what library redirection, if any, you will be using for the remote journals and associated journal receivers.
- 3. Ensure that all selected libraries exist on the target systems. You must consider whether or not library redirection will be used when adding the remote journal.
- Create the appropriate local journal if it does not already exist.
- 5. Configure and activate the communications protocol you have chosen to use.
- 6. After you have configured the communications protocol, it must be active while you are using the remote journal function.
 - For example, if you are using the OptiConnect for IBM i bus transport method, then the OptiConnect for IBM i subsystem, QSOC, must be active. QSOC must be active for both the source system and the target system, and the appropriate controllers and devices must be varied on. If you are using a SNA communications transport, vary on the appropriate line, controller, and devices and ensure subsystem QCMN is active on both systems. Start of change If you are using TCP/IP or Sockets IPv6, you must start TCP/IP by using the Start TCP/IP (STRTCP) command, including the distributed data management (DDM) servers. If you are using data port, you must configure a cluster, make sure that the cluster is active, and start the internet Daemon (inetd) server using the Start TCP/IP Server (STRTCPSVR) command.End of change
- 7. If one does not already exist, create the appropriate relational database (RDB) directory entry that will be used to define the communications protocol for the remote journal environment. When TCP communications are being used to connect to an independent

disk pool, the RDB entry to the independent disk pool must have the Relational database value set to the target system's local RDB entry and the relational database alias value set to the independent disk pool's name.

8. Now you should be able to see the remote database connection by issuing the WRKRDBDIRE command.

```
Work with Relational Database Directory Entries

Position to . . . .

Type options, press Enter.

1=Add 2=Change 4=Remove 5=Display details 6=Print details

Remote

Option Entry Location Text

SYS1 system1
SYS2 system2
MYSYSTEM *LOCAL Entry added by system
```

Bottom

F3=Exit F5=Refresh F6=Print list F12=Cancel F22=Display entire field (C) COPYRIGHT IBM CORP. 1980, 2007.

Adding a Remote Journal

Adding a remote journal creates a remote journal on a target system or independent disk pool and associates that remote journal with the journal on the source system. This occurs if this is the first time the remote journal is being established for a journal. The journal on the source system can be either a local or remote journal.

If a remote journal environment has previously been established, adding a remote journal reassociates the remote journal on the target system with the journal on the source system.

You can establish and associate a remote journal on a target system with a journal on the source system by one of the following methods:

- System i Navigator.
- Add the Remote Journal (QjoAddRemoteJournal) API on the source system.
- Add the Remote Journal (ADDRMTJRN) command on the source system.
- What Happens During Add Remote Journal Processing?
- Guidelines For Adding a Remote Journal

What Happens During Add Remote Journal Processing?

The processing that takes place as part of adding a remote journal includes the following:

• A check is performed on the target system to verify that the user profile adding the remote journal exists. A user profile with the same name as the user profile which is adding a

- remote journal must exist on the target system. If the profile does not exist on the target system, then an exception is signaled, and the processing ends.
- A check is performed to verify that the target system has a library by the same name as
 the library for the journal on the source system. If the library does not exist on the target
 system, then an exception is signaled, and the processing ends.
- A check is performed on the target system to determine if a journal by the same qualified name as the journal on the source system already exists. If a journal already exists, it can be used for the remainder of the add remote journal processing if it meets the following criteria:
 - It is a remote journal.
 - It was previously associated with this same source journal or part of the same remote journal network.
 - 3. The type of the remote journal matches the specified remote journal type.
- If a journal was found, but does not meet the preceding criteria, then an exception is signaled, and the processing ends. Otherwise, the remote journal is used for the rest of the add remote journal processing.
- If no journal is found on the specified target system, then a remote journal is created on the target system. The new remote journal has the same configuration, authority, and audit characteristics of the source journal. The journal that is created has a journal type of *REMOTE.

When adding the remote journal, you must specify the type of remote journal to add. The remote journal type influences the library redirection rules and other operational characteristics for the journal.

Guidelines For Adding a Remote Journal

You should observe the following guidelines for adding a remote journal:

- You can only associate a remote journal with a single source journal.
 - Note: The same remote journal can then have additional remote journals that are associated with it that are located on other target systems. This is the cascade configuration that is shown in Network configurations for remote journals.
- The remote journal will only have its attached receiver populated with journal entries that are replicated from the corresponding journal receiver on the source system. No journal entries can be directly deposited to a remote journal.
- A maximum of 255 remote journals can be associated with a single journal on a source system. This can be any combination of asynchronously maintained or synchronously maintained remote journals.

To Add a Remote Journal

The following is an example using the physical file QGPL/TESTPF setup to have remote journaling enabled to a second system.

1. Create the physical file.:

```
> CRTPF FILE(QGPL/TESTPF) RCDLEN(10)
File TESTPF created in library QGPL.
Member TESTPF added to file TESTPF in QGPL.
```

Create the local journal receiver and journals, and enable the journaling of the physical file created:



- > crtjrnrcv jrnrcv(qgpl/jrcvrmt)
 Journal receiver JRCVRMT created in library QGPL
- > crtjrn jrn(qgpl/jrnrmt) jrnrcv(qgpl/jrcvrmt) fixlendta(*job *usr *pgm *sysseq)
 Journal JRNRMT created in library QGPL

strjrnpf file(qgpl/testpf) jrn(qgpl/testpf)
1 of 1 files have started journaling

3. Add the remote journal:

> addrmtjrn rdb(sys2) srcjrn(qgpl/JRNRMT) rmtjrntype(*TYPE2)
Remote journal JRNRMT in QGPL was added

4. Activate the remote journaling:

> chgrmtjrn rbd(sys2) srcjrn(qgpl/jrnrmt) jrnstate(*active)
Remote journal JRNRMT in library QGPL was activated



Part IV

Using Oracle GoldenGate for DB2 for z/OS

With Oracle GoldenGate for DB2 for z/OS, you can perform initial loads and capture transactional data from supported DB2 for z/OS versions and replicate the data to a DB2 for z/OS database or other supported Oracle GoldenGate targets, such as an Oracle Database.

Oracle GoldenGate for DB2 for z/OS is installed and runs remotely on Linux, zLinux, or AIX.

Oracle GoldenGate for DB2 for z/OS supports data filtering, mapping, and transformations unless noted otherwise in this documentation.

Topics:

- Understanding What's Supported for DB2 for z/OS
 This chapter contains information on database and table features supported byOracle GoldenGate for DB2 z/OS.
- Preparing the DB2 for z/OS Database for Oracle GoldenGate
- Preparing the DB2 for z/OS Transaction Logs for Oracle GoldenGate



9

Understanding What's Supported for DB2 for z/OS

This chapter contains information on database and table features supported by Oracle Golden Gate for DB2 z/OS.

Topics:

- Supported DB2 for z/OS Data Types
- Non-Supported DB2 for z/OS Data Types
- Supported Objects and Operations for DB2 z/OS
- Non-Supported Objects and Operations for DB2 for z/OS

Supported DB2 for z/OS Data Types

This section lists the DB2 for z/OS data types that Oracle GoldenGate supports and any limitations of this support.

- Oracle GoldenGate does not perform character set conversion for columns that could contain multi-byte data. This includes <code>GRAPHIC</code>, <code>VARGRAPHIC</code> and <code>DBCLOB</code> data types, as well as <code>CHAR</code>, <code>VARCHAR</code>, and <code>CLOB</code> for tables defined with <code>ENCODING_SCHEME</code> of 'M' (multiple CCSID set or multiple encoding schemes) or 'U' (Unicode). Such data is only supported if the source and target systems are the same CCSID.
- Oracle GoldenGate supports ASCII, EBCDIC, and Unicode data format. Oracle GoldenGate converts between ASCII and EBCDIC data automatically. Unicode is not converted.
- Oracle GoldenGate supports most DB2 data types except those listed in Non-Supported DB2 for z/OS Data Types.

Limitations of Support

- The support of range and precision for floating-point numbers depends on the host machine. In general, the precision is accurate to 16 significant digits, but you should review the database documentation to determine the expected approximations. Oracle GoldenGate rounds or truncates values that exceed the supported precision.
- Oracle GoldenGate does not support the filtering, column mapping, or manipulation of large objects greater than 4K in size. You can use the full Oracle GoldenGate functionality for objects that are 4K or smaller.
- Oracle GoldenGate supports the default TIMESTAMP and the TIMESTAMP with TIMEZONE to up to 9 digit fractional value, but no further.

Non-Supported DB2 for z/OS Data Types

This section lists DB2 for z/OS data types that Oracle GoldenGate does not support. Data that is not supported may affect the integrity of the target data in relation to the source data.

- XML
- User-defined types
- Negative dates

Supported Objects and Operations for DB2 z/OS

This section lists the database objects and types of operations that Oracle GoldenGate supports.

- Extraction and replication of DML operations on DB2 for z/OS tables that contain rows of up to 512KB in length. This size exceeds the maximum row size of DB2.
- INSERT operations from the IBM LOAD utility are supported for change capture if the utility is run with LOG YES and SHRLEVEL CHANGE, and the source tables that are being loaded have DATA CAPTURE CHANGES enabled (required by Oracle GoldenGate) and are specified in the Oracle GoldenGate Extract configuration. Oracle GoldenGate also supports initial loads with the LOAD utility to instantiate target tables during initial synchronization.
- Oracle GoldenGate supports the maximum number of columns per table, which is supported by the database.
- Oracle GoldenGate supports the maximum column size that is supported by the database.
- Extraction and replication of data that is stored using DB2 data compression (CREATE TABLESPACE COMPRESS YES).
- Capture from temporal history tables is supported.
- TRUNCATE TABLE is supported, but because this command issues row deletes to perform the truncate, they are shown in Oracle GoldenGate statistics as such, and not as a truncate operation. To replicate a TRUNCATE, the Replicat process uses a DELETE operation without a WHERE clause.
- TRUNCATES are always captured from a DB2 for z/OS source, but can be ignored by Replicat if the IGNORETRUNCATES parameter is used in the Replicat parameter file.
- UNICODE columns in EBCDIC tables are supported.
- Supported options with SHOWTRANS

```
SHOWTRANS [transaction_ID] [COUNT n]
[DURATION duration unit]
[FILE file name] |
```

transaction ID and count cannot be specified together.

transaction ID and duration cannot be specified together.

Options supported with SKIPTRANS and FORCETRANS:

```
SKIPTRANS transaction_ID
[FORCE] FORCETRANS transaction ID [FORCE]
```

Non-Supported Objects and Operations for DB2 for z/OS

The following objects and operations are not supported by Oracle GoldenGate on DB2 for z/OS:

- Extraction or replication of DDL operations
- Clone tables
- Data manipulation, including compression, that is performed within user-supplied DB2 exit routines, such as:
 - Date and time routines
 - Edit routines (CREATE TABLE EDITPROC)
 - Validation routines (CREATE TABLE VALIDPROC)
- Replicating with BATCHSQL is not fully functional for DB2 for z/OS. Non-insert operations are
 not supported so any update or delete operations will cause Replicat to drop temporarily
 out of BATCHSQL mode. The transactions will stop and errors will occur.



10

Preparing the DB2 for z/OS Database for Oracle GoldenGate

Learn how to prepare your database and environment to support Oracle GoldenGate. **Topics:**

- Preparing Tables for Processing
- Configuring a Database Connection
- Monitoring Processes
- Supporting Globalization Functions

Preparing Tables for Processing

You must perform the following tasks to prepare your tables for use in an Oracle GoldenGate environment.

- Disabling Triggers and Cascade Constraints
- Assigning Row Identifiers
- Handling ROWID Columns

Disabling Triggers and Cascade Constraints

Disable triggers, cascade delete constraints, and cascade update constraints on the target tables, or alter them to ignore changes made by the Oracle GoldenGate database user. Oracle GoldenGate replicates DML that results from a trigger or cascade constraint. If the same trigger or constraint gets activated on the target table, it becomes redundant because of the replicated version, and the database returns an error. Consider the following example, where the source tables are <code>emp_src</code> and <code>salary_src</code> and the target tables are <code>emp_targ</code> and <code>salary_targ</code>.

- A delete is issued for emp src.
- It cascades a delete to salary src.
- Oracle GoldenGate sends both deletes to the target.
- The parent delete arrives first and is applied to emp targ.
- The parent delete cascades a delete to salary targ.
- The cascaded delete from salary src is applied to salary targ.
- The row cannot be located because it was already deleted in step 5.

Assigning Row Identifiers

Oracle GoldenGate requires some form of unique row identifier on the source and target tables to locate the correct target rows for replicated updates and deletes.

- How Oracle GoldenGate Determines the Kind of Row Identifier to Use
- Using KEYCOLS to Specify a Custom Key

How Oracle GoldenGate Determines the Kind of Row Identifier to Use

Unless a KEYCOLS clause is used in the TABLE or MAP statement, Oracle GoldenGate selects a row identifier to use in the following order of priority:

- 1. Primary key
- 2. First unique key alphanumerically that does not contain a timestamp or non-materialized computed column.
- 3. If none of the preceding key types exist (even though there might be other types of keys defined on the table) Oracle GoldenGate constructs a pseudo key of all columns that the database allows to be used in a unique key, excluding those that are not supported by Oracle GoldenGate in a key or those that are excluded from the Oracle GoldenGate configuration.

Note:

If there are other, non-usable keys on a table or if there are no keys at all on the table, Oracle GoldenGate logs an appropriate message to the report file. Constructing a key from all of the columns impedes the performance of Oracle GoldenGate on the source system. On the target, this key causes Replicat to use a larger, less efficient WHERE clause.

Using **KEYCOLS** to Specify a Custom Key

If a table does not have one of the preceding types of row identifiers, or if you prefer those identifiers not to be used, you can define a substitute key if the table has columns that always contain unique values. You define this substitute key by including a KEYCOLS clause within the Extract TABLE parameter and the Replicat MAP parameter. The specified key will override any existing primary or unique key that Oracle GoldenGate finds. For more information, see Reference for Oracle GoldenGate.

Handling ROWID Columns

Any attempt to insert into a target table that includes a column with a data type of ROWID GENERATED ALWAYS (the default) will fail with the following ODBC error:

ODBC error: SQLSTATE 428C9 native database error -798. {DB2 FOR OS/390}{ODBC DRIVER} {DSN08015} DSNT408I SQLCODE = -798, ERROR: YOU CANNOT INSERT A VALUE INTO A COLUMN THAT IS DEFINED WITH THE OPTION GENERATED ALWAYS. COLUMN NAME ROWIDCOL.

You can do one of the following to prepare tables with ROWID columns to be processed by Oracle GoldenGate:

- Ensure that any ROWID columns in target tables are defined as GENERATED BY DEFAULT.
- If it is not possible to change the table definition, you can work around it with the following procedure.



To Work Around ROWID GENERATE ALWAYS:

1. For the source table, create an Extract TABLE statement, and use a COLSEXCEPT clause in that statement that excludes the ROWID column. For example:

```
TABLE tab1, COLSEXCEPT (rowidcol);
```

The COLSEXCEPT clause excludes the ROWID column from being captured and replicated to the target table.

- 2. For the target table, ensure that Replicat does not attempt to use the ROWID column as the key. This can be done in one of the following ways:
 - Specify a primary key in the target table definition.
 - If a key cannot be created, create a Replicat MAP parameter for the table, and use a KEYCOLS clause in that statement that contains any unique columns except for the ROWID column. Replicat will use those columns as a key. For example:

```
MAP tab1, TARGET tab1, KEYCOLS (num, ckey);
```

For more information about KEYCOLS, see Assigning Row Identifiers.

Configuring a Database Connection

This section contains instructions for setting up the Extract and Replicat connections to a SQL Server database.

- Database Configuration for DB2 z/OS
- Database User for Oracle GoldenGate Processes
- Setting Initialization Parameters
- · Specifying the Path to the Initialization File
- Ensuring ODBC Connection Compatibility
- Specifying the Number of Connection Threads

Database Configuration for DB2 z/OS

No special DB2 z/OS database settings are required for Oracle GoldenGate.

Database User for Oracle GoldenGate Processes

Oracle GoldenGate requires a database user account. Create this account and assign privileges according to the following guidelines.

Assign the DB2 privileges listed in Table 10-1 to the user by which Extract and Replicat will be running. These are in addition to any permissions that DB2 ODBC requires. All Extract privileges apply to initial-load and log-based Extract processes, except where noted. The following authorities can be provided by granting either SYSCTRL or DBADM plus SQLADM authority to the user running the Oracle GoldenGate processes.



Table 10-1 Privileges Needed by Oracle GoldenGate for DB2 z/OS

User privilege	Extract	Replicat
MONITOR2	X	
(does not apply to initial-load Extract)		
SELECT ON the following SYSIBM tables:	X	Х
SYSTABLES		
SYSCOLUMNS		
SYSTABLEPART		
SYSKEYS		
SYSINDEXES		
SYSCOLAUTH		
SYSDATABASE		
SYSFOREIGNKEYS		
SYSPARMS		
SYSRELS		
SYSROUTINES		
SYSSYNONYMS		
SYSTABAUTH		
SYSAUXRELS		
SELECT on source tables ¹	X	
INSERT, UPDATE, DELETE on target tables		Х
CREATE TABLE ²		Χ
EXECUTE on ODBC plan (default is DSNACLI)	Х	
Privileges required by ${\tt SQLEXEC}$ procedures or queries that you will be using. $^{\tt 3}$	X	Х

¹ SELECT on source tables required only if tables contain LOB columns, or for an initial-load Extract, if used.

Setting Initialization Parameters

The following DB2 for z/OS initialization parameters apply to Oracle GoldenGate and must be set correctly before starting Oracle GoldenGate processes.

- MVSDEFAULTSSID: set to the DB2 subsystem.
- LOCATION: set to the DB2 location name as stored in the DB2 Boot Strap Dataset.
- MVSATTACHTYPE: set to RRSAF (Recoverable Resource Manager Services Attachment Facility) or CAF (Call Attachment Facility). IBM recommends using RRSAF.
- MULTICONTEXT: set to 1 if using RRSAF.
- PLANNAME: set to the DB2 plan. The default plan name is DSNACLI.

Do not use the CURRENTAPPENSCH initialization parameter (keyword).



² Required if using ADD CHECKPOINTTABLE in GGSCI to use the database checkpoint feature.

³ SQLEXEC enables stored procedures and queries to be executed by an Oracle GoldenGate process.



When using the CAF attachment type, you must use the Oracle GoldenGate DBOPTIONS parameter with the NOCATALOGCONNECT option in the parameter file of any Extract or Replicat process that connects to DB2. This parameter disables the usual attempt by Oracle GoldenGate to obtain a second thread for the DB2 catalog. Otherwise, you will receive error messages, such as: ODBC operation failed: Couldn't connect to data source for catalog gueries.

Specifying the Path to the Initialization File

Specify the ODBC initialization file by setting the DSNAOINI environment variable in the z/OS UNIX profile, as in the following example:

export DSNAOINI="/etc/odbc810.ini"

Ensuring ODBC Connection Compatibility

To ensure that you configure the DB2 ODBC initialization file correctly, follow the guidelines in the DB2 UDB for z/OS ODBC Guide and Reference manual. One important consideration is the coding of the open and close square brackets (the [character and the] character). The square bracket characters are "variant" characters that are encoded differently in different coded character set identifiers (CCSID), but must be of the IBM-1047 CCSID in the ODBC initialization file. DB2 ODBC does not recognize brackets of any other CCSID. Note the following:

- The first (or open) bracket must use the hexadecimal characters X'AD' (0xAD).
- The second (or close) bracket must use the hexadecimal characters X'BD' (0xBD).

To set the correct code for square brackets, use any of the following methods.

- Use the hex command in OEDIT and change the hex code for each character appropriately.
- Use the iconv utility to convert the ODBC initialization file. For example, to convert from CCSID IBM-037 to IBM-1047, use the following command:

```
iconv -f IBM-037 -t IBM-1047 ODBC.ini > ODBC-1047.ini
mv ODBC-1047.ini ODBC.ini
```

• Change your terminal emulator or terminal configuration to use CCSID IBM-1047 when you create or alter the file.

Specifying the Number of Connection Threads

Every Oracle GoldenGate process makes a database connection. Depending on the number of processes that you will be using and the number of other DB2 connections that you expect, you might need to adjust the following DB2 system parameters on the DSNTIPE DB2 Thread Management Panel:

- MAX USERS (macro DSN6SYSP CTHREAD)
- MAX TSO CONNECT (macro DSN6SYSP IDFORE)
- MAX BATCH CONNECT (macro DSN6SYSP IDBACK)



If using RRSAF, allow:

- Two DB2 threads per process for each of the following:
 - Extract
 - Replicat
 - The GGSCI command DBLOGIN (logs into the database)
 - DEFGEN utility (generates data definitions for column mapping)
- One extra DB2 thread for Extract for IFI calls.
- One extra DB2 thread for each SQLEXEC parameter statement that will be issued by each
 Extract and Replicat process. For more information about SQLEXEC, see the Reference for
 Oracle GoldenGate.

If using CAF, there can be only one thread per Oracle GoldenGate process.

Monitoring Processes

These sections provide information about monitoring Oracle GoldenGate with z/OS system facilities.

Interpreting Statistics for Update Operations

Interpreting Statistics for Update Operations

The actual number of DML operations that are executed on the DB2 database might not match the number of extracted DML operations that are reported by Oracle GoldenGate. DB2 does not log update statements if they do not physically change a row, so Oracle GoldenGate cannot detect them or include them in statistics.

Supporting Globalization Functions

Oracle GoldenGate provides globalization support and you should take into consideration when using this support.

- Replicating From a Source that Contains Both ASCII and EBCDIC
- Specifying Multi-Byte Characters in Object Names

Replicating From a Source that Contains Both ASCII and EBCDIC

When replicating to or from a DB2 source system to a target that has a different character set, some consideration must be given to the encoding of the character data on the DB2 source if it contains a mix of ASCII and EBCDIC data. Character set conversion by any given Replicat requires source data to be in a single character set.

The source character set is specified in the trail header. Thus, the Oracle GoldenGate trail can contain either ASCII or EBCDIC data, but not both. Unicode tables are processed without any special configuration and are exempt from the one-character set requirement.

With respect to a source that contains both character encoding types, you have the following options:

 You can use one Extract for all of your tables, and have it write the character data to the trail as either ASCII or as EBCDIC.



 You can use different Extracts: one Extract to write the ASCII character data to a trail, and another Extract to write the EBCDIC character data to a different trail. You then associate each trail with its own data pump process and Replicat process, so that the two data streams are processed separately.

To output the correct character set in either of those scenarios, use the TRAILCHARSETASCII and TRAILCHARSETEBCDIC parameters. The default is TRAILCHARSETEBCDIC. Without these parameters, ASCII and EBCDIC data are written to the trail as-is. When using these parameters, note the following:

- If used on a single-byte DB2 subsystem, these parameters cause Extract to convert all of
 the character data to either the ASCII or EBCDIC single-byte CCSID of the subsystem to
 which Extract is connected, depending on which parameter is used (except for Unicode,
 which is processed as-is).
- If used on a multi-byte DB2 subsystem, these parameters cause Extract to capture only ASCII or EBCDIC tables (and Unicode). Character data is written in either the ASCII or EBCDIC mixed CCSID (depending on the parameter used) of the DB2 z/OS subsystem to which Extract is connected.

Specifying Multi-Byte Characters in Object Names

If the name of a schema, table, column, or stored procedure in a parameter file contains a multi-byte character, the name must be double-quoted. For more information about specifying object names, see *Administering Oracle GoldenGate*.



11

Preparing the DB2 for z/OS Transaction Logs for Oracle GoldenGate

Learn how to configure the DB2 transaction logging to support data capture by Oracle GoldenGate Extract.

Topics:

Making Transaction Data Available

Making Transaction Data Available

Oracle GoldenGate can extract DB2 transaction data from the active and archived logs. Follow these guidelines to configure the logs so that Extract can capture data.

- Enabling Change Capture
- Enabling Access to Log Records
- · Sizing and Retaining the Logs
- Using Archive Logs on Tape
- Controlling Log Flushes

Enabling Change Capture

Follow these steps to configure DB2 to log data changes in the expanded format that is supplied by the DATA CAPTURE CHANGES feature of the CREATE TABLE and ALTER TABLE commands. This format provides Oracle GoldenGate with the entire before and after images of rows that are changed with update statements.

- 1. From the Oracle GoldenGate directory, run GGSCI.
- 2. Log on to DB2 from GGSCI as a user that has ALTER TABLE privileges.

```
DBLOGIN SOURCEDB DSN, USERID user[, PASSWORD password][, encryption options]
```

3. Issue the following command. where *table* is the fully qualified name of the table. You can use a wildcard to specify multiple table names but not owner names.

```
ADD TRANDATA table
```

By default, ADD TRANDATA issues the following command:

ALTER TABLE name DATA CAPTURE CHANGES;

Enabling Access to Log Records

Activate DB2 Monitor Trace Class 1 ("TRACE (MONITOR) CLASS (1)") so that DB2 allows Extract to read the active log. The default destination of OPX is sufficient, because Oracle GoldenGate does not use a destination.

To Start the Trace Manually

- 1. Log on to DB2 as a DB2 user who has the TRACE privilege or at least SYSOPR authority.
- 2. Issue the following command:

```
start trace(monitor) class(1) scope(group)
```

To Start the Trace Automatically When DB2 is Started

Do either of the following:

- Set MONITOR TRACE to "YES" on the DSNTIPN installation tracing panel.
- Set 'DSN6SYSP MON=YES' in the DSNTIJUZ installation job, as described in the DB2 UDB Installation Guide.



The primary authorization ID, or one of the secondary authorization IDs, of the ODBC plan executor also must have the MONITOR2 privilege.

Sizing and Retaining the Logs

When tables are defined with DATA CAPTURE CHANGES, more data is logged than when they are defined with DATA CAPTURE NONE. If any of the following is true, you might need to increase the number and size of the active and archived logs.

- Your applications generate large amounts of DB2 data.
- Your applications have infrequent commits.
- You expect to stop Extract for long periods of time.
- Your network is unreliable or slow.

To control log retention, use the DSN6LOGP MAXARCH system parameter in the DSNTIJUZ installation job.

Retain enough log data so that Extract can start again from its checkpoints after you stop it or after an unplanned outage. Extract must have access to the log that contains the start of the oldest uncommitted unit of work, and all logs thereafter.

If data that Extract needs during processing was not retained, either in online or archived logs, one of the following corrective actions might be required:

- Alter Extract to capture from a later point in time for which log data is available (and accept possible data loss on the target).
- Resynchronize the source and target tables, and then start the Oracle GoldenGate environment over again.



Note:

The IBM documentation makes recommendations for improving the performance of log reads. In particular, you can use large log output buffers, large active logs, and make archives to disk.

Using Archive Logs on Tape

Oracle GoldenGate can read DB2 archive logs on tape, but it will degrade performance. For example, DB2 reserves taped archives for a single recovery task. Therefore, Extract would not be able to read an archive tape that is being used to recover a table until the recovery is finished. You could use DFHSM or an equivalent tools to move the archive logs in a seamless manner between online DASD storage and tape, but Extract will have to wait until the transfer is finished. Delays in Extract processing increase the latency between source and target data.

Controlling Log Flushes

When reading the transaction log, Extract does not process a transaction until it captures the commit record. If the commit record is on a data block that is not full, it cannot be captured until more log activity is generated to complete the block. The API that is used by Extract to read the logs only retrieves full physical data blocks.

A delay in receiving blocks that contain commits can cause latency between the source and target data. If the applications are not generating enough log records to fill a block, Extract generates its own log records by issuing SAVEPOINT and COMMIT statements, until the block fills up one way or the other and is released.

In a data sharing group, each API call causes DB2 to flush the data blocks of all active members, eliminating the need for Extract to perform flushes.

To prevent Extract from performing flushes, use the Extract parameter TRANLOGOPTIONS with the NOFLUSH option.



Part V

Using Oracle GoldenGate for MySQL

With Oracle GoldenGate for MySQL, you can perform initial loads and capture transactional data and table changes from supported MySQL versions and replicate the data and table changes to a MySQL database or replicate the data to other supported Oracle GoldenGate targets, such as an Oracle Database.

Oracle GoldenGate for MySQL supports data filtering, mapping, and transformations unless noted otherwise in this documentation.

This part describes tasks for configuring and running Oracle GoldenGate for MySQL and supported variants, such as MariaDB, Amazon RDS for MySQL, and Amazon Aurora MySQL.

Topics:

- Understanding What's Supported for MySQL
 This chapter contains information on database and table features supported by Oracle GoldenGate.
- Preparing and Configuring the System for Oracle GoldenGate
- Using DDL Replication



Understanding What's Supported for MySQL

This chapter contains information on database and table features supported by Oracle GoldenGate.

Topics:

- Character Sets in MySQL
- Supported MySQL Data Types
- Non-Supported MySQL Data Types
- Supported Objects and Operations for MySQL
- Non-Supported Objects and Operations
- System Schemas

Character Sets in MySQL

MySQL provides a facility that allows users to specify different character sets at different levels.

Level	Example
Database	create database test charset utf8;
Table	create table test(id int, name char(100)) charset utf8;
Column	<pre>create table test (id int, name1 char(100) charset gbk, name2 char(100) charset utf8));</pre>

Limitations of Support

- When you specify the character set of your database as utf8mb4/utf8, the default collation is utf8mb4_unicode_ci/utf8_general_ci. If you specify collation_server=utf8mb4_bin, the database interprets the data as binary. For example, specifying the CHAR column length as four means that the byte length returned is 16 (for utf8mb4) though when you try to insert data more than four bytes the target database warns that the data is too long. This is the limitation of database so Oracle GoldenGate does not support binary collation. To overcome this issue, specify collation_server=utf8mb4_bin when the character set is utf8mb4 and collation_server=utf8_bin for utf8.
- The following character sets are not supported:

```
armscii8
keybcs2
utf16le
geostd8
```

Supported MySQL Data Types

MySQL supports the following data types:

- CHAR
- VARCHAR
- INT
- TINYINT
- SMALL INT
- MEDIUM INT
- BIG INT
- DECIMAL
- FLOAT
- DOUBLE
- DATE
- YEAR
- DATETIME

TIME

- TIMESTAMP
- BINARY
- VARBINARY
- TEXT
- TINYTEXT
- MEDIUMTEXT
- LONGTEXT
- BLOB
- TINYBLOB
- MEDIUMBLOB
- LONGBLOB
- ENUM
- BIT (M)
- Limitations and Clarifications

Limitations and Clarifications

When running Oracle GoldenGate for MySQL, be aware of the following:

- Functional indexes are not supported for Capture or Delivery.
- Oracle GoldenGate does not support BLOB or TEXT types when used as a primary key.



- Oracle GoldenGate supports UTF8 and UCS2 character sets. UTF8 data is converted to UTF16 by Oracle GoldenGate before writing it to the trail.
- UTF32 is not supported by Oracle GoldenGate.
- Oracle GoldenGate supports a TIME type range from 00:00:00 to 23:59:59.
- Oracle GoldenGate supports timestamp data from 0001/01/03:00:00:00 to 9999/12/31:23:59:59. If a timestamp is converted from GMT to local time, these limits also apply to the resulting timestamp. Depending on the time zone, conversion may add or subtract hours, which can cause the timestamp to exceed the lower or upper supported limit.
- When the time zone of the Oracle GoldenGate installation server does not match the time zone of the source database server, then the TIMESTAMP data sent to the target database will differ from the source database.
 - For Classic Architecture, create a session variable for Oracle GoldenGate, called ${\tt TZ}$, and set it equal to the time zone value of the database.
- Oracle GoldenGate does not support negative dates.
- The support of range and precision for floating-point numbers depends on the host machine. In general, the precision is accurate to 16 significant digits, but you should review the database documentation to determine the expected approximations. Oracle GoldenGate rounds or truncates values that exceed the supported precision.
- When you use ENUM type in non-strict sql_mode, the non-strict sql_mode does not prevent you from entering an invalid ENUM value and an error will be returned. To avoid this situation, do one of the following:
 - Use sql_mode as STRICT and restart Extract. This prevents users from entering invalid values for any of the data types. An IE user can only enter valid values for those data types.
 - Continue using non-strict sql mode, but do not use ENUM data types.
 - Continue using non-strict sql_mode and use ENUM data types with valid values in the database. If you specify invalid values, the database will silently accept them and Extract will abend.
- To preserve transaction boundaries for a MySQL target, create or alter the target tables to the InnoDB transactional database engine instead of the MyISAM engine. MyISAM will cause Replicat records to be applied as they are received, which does not guarantee transaction integrity even with auto-commit turned off. You cannot roll back a transaction with MyISAM.
- Extraction and replication from and to views is not supported.
- Transactions applied by the slave are logged into the relay logs and not into the slave's binlog. If you want a slave to write transactions the binlog that it receives from the master, you need to start the replication slave with the log slave-updates option as 1 in my.cnf. This is in addition to the other binary logging parameters. After the master's transactions are in the slave's binlog, you can then setup a regular capture on the slave to capture and process the slave's binlog.

Non-Supported MySQL Data Types

Oracle GoldenGate for MySQL does not support the following data types:

All spatial types (Geometry and so on), JSON, SET, XML



Note:

Extract abends if it is configured to capture from tables that contain any of the unsupported data types, so ensure that Extract is not configured to capture from tables containing columns of unsupported data types.

Supported Objects and Operations for MySQL

Oracle GoldenGate for MySQL supports the following objects and operations:

- Oracle GoldenGate supports the following DML operations on source and target database transactional tables:
 - Insert operation
 - Update operation (compressed included)
 - Delete operation (compressed included)
 - Truncate operation
- Oracle GoldenGate supports the extraction and replication of DDL (data definition language) operations.
- Oracle GoldenGate supports transactional tables up to the full row size and maximum number of columns that are supported by MySQL and the database storage engine that is being used. InnoDB supports up to 1017 columns.
- Oracle GoldenGate supports the AUTO_INCREMENT column attribute. The increment value is
 captured from the binary log by Extract and applied to the target table in a Replicat insert
 operation.
- Oracle GoldenGate can operate concurrently with MySQL native replication.
- Oracle GoldenGate supports the DYNSQL feature for MySQL.

Note:

XA transactions are not supported for capture and any XA transactions logged in binlog cause Extract to abend. So, you must not use XA transactions against a database that Extract is configured to capture.

If XA transactions are being used for databases that are not configured for Oracle GoldenGate capture, then exclude those databases from logging into MySQL binary logs by using the parameter <code>binlog-ignore-db</code> in the MySQL server configuration file.

Limitations on Automatic Heartbeat Table support are as follows:

- Ensure that the database in which the heartbeat table is to be created already exists to avoid errors when adding the heartbeat table.
- In the heartbeat history lag view, the information in fields like heartbeat_received_ts, incoming_heartbeat_age, and outgoing_heartbeat_age are shown with respect to the system time. You should ensure that the operating system time is setup with the correct and current time zone information.



Non-Supported Objects and Operations

Oracle GoldenGate for MySQL does not support the following objects and operations:

- Invisible columns
- The Oracle GoldenGate BATCHSQL feature.
- Array fetching during initial load.
- The following character sets are not supported:

```
ULIB_CS_ARMSCII8, /* American National 166-9 */
ULIB_CS_GEOSTD8, /* Geogian Standard */
ULIB_CS_KEYBCS2, /* Kemennicky MS-DOS
```

- Capturing NLS LOB data using the FETCHMOCOLS and FETCHMODCOLEXCEPT TABLE options is not supported when DDL is enabled.
- Renaming tables.
- DDL statements inside stored procedures is not supported.
- When the time zone of the Oracle GoldenGate installation server does not match the time zone of the source database server, then the TIMESTAMP data sent to the target database will differ from the source database. For Oracle GoldenGate Microservices installations, regardless of the time zones being the same, Extract will resolve the time zone to UTC.
 Determine the source database time zone by running the following query:

```
select @@system time zone;
```

This will return a time zone value, such as PDT.

Create a session variable for Oracle GoldenGate, called ${\tt TZ}$, and set it equal to the time zone value of the database.

- Extraction and replication from and to views is not supported.
- Transactions applied by the slave are logged into the relay logs and not into the slave's binlog. If you want a slave to write transactions the binlog that it receives from the master, you need to start the replication slave with the log slave-updates option as 1 in my.cnf. This is in addition to the other binary logging parameters. After the master's transactions are in the slave's binlog, you can then setup a regular capture on the slave to capture and process the slave's binlog.

System Schemas

The following schemas or objects are not automatically replicated by Oracle GoldenGate unless they are explicitly specified without a wildcard.

- 'information schema'
- 'performance schema'
- 'mysql'
- 'sys'



Preparing and Configuring the System for Oracle GoldenGate

Learn about how to prepare the system for running Oracle GoldenGate and how to configure it with your MySQL database.

Topics:

- Database User for Oracle GoldenGate Processes for MySQL
- Ensuring Data Availability
- Setting Logging Parameters
- Adding Host Names
- Setting the Session Character Set
- Preparing Tables for Processing
- Changing the Log-Bin Location
- Configuring Bi-Directional Replication
- Configuring MySQL for Remote Capture
- Configuring a Two-way SSL Connection in MySQL Capture and Delivery
- Capturing using a MySQL Replication Slave
- Other Oracle GoldenGate Parameters for MySQL
- Positioning Extract to a Specific Start Point

Database User for Oracle GoldenGate Processes for MySQL

Requirements for the database user for Oracle GoldenGate processes are as follows:

- Create a database user that is dedicated to Oracle GoldenGate. It can be the same user for all of the Oracle GoldenGate processes that must connect to a database:
 - Extract (source database)
 - Replicat (target database)
 - DEFGEN (source or target database)
- To use DDL the MySQL user must have privileges to install the database plug-ins.
- To preserve the security of your data, and to monitor Oracle GoldenGate processing accurately, do not permit other users, applications, or processes to log on as, or operate as, the Oracle GoldenGate database user.
- Keep a record of the database users. They must be specified in the Oracle GoldenGate parameter files with the USERID parameter.
- The Oracle GoldenGate user requires read access to the INFORMATION SCHEMA database.
- The Oracle GoldenGate user requires the following user privileges.



Privilege	Source Extract	Target Replicat	Purpose
SELECT	Yes	Yes	Connect to the database and select object definitions
REPLICATION SLAVE	Yes	NA	Connect and receive updates from the replication master's binary log.
CREATE CREATE VIEW EVENT INSERT UPDATE DELETE	Yes	Yes	Source and target database heartbeat and checkpoint table creation, and data record generation and purging
DROP	Yes	Yes	Dropping a Replicat checkpoint table or deleting a heartbeat table implementati on
EXECUTE	Yes	Yes	To execute stored procedures.
INSERT, UPDATE, DELETE on target tables	NA	Yes	Apply replicated DML to target objects
DDL privileges on target objects (if using DDL support)	NA	Yes	Issue replicated DDL on target objects

- To capture binary log events, an Administrator must provide the following privileges to the Extract user:
 - Read and Execute permissions for the directory where the MySQL configuration file (my.cnf) is located
 - Read permission for the MySQL configuration file (my.cnf)
 - Read and Execute permissions for the directory where the binary logs are located
 - Read and Execute permission for the tmp directory



Ensuring Data Availability

Retain enough binary log data so that if you stop Extract or there is an unplanned outage, Extract can start again from its checkpoints. Extract must have access to the binary log that contains the start of the oldest uncommitted unit of work, and all binary logs thereafter. The recommended retention period is at least 24 hours worth of transaction data, including both active and archived information. You might need to do some testing to determine the best retention time given your data volume and business requirements.

If data that Extract needs during processing was not retained, either in active or backup logs, one of the following corrective actions might be required:

- Alter Extract to capture from a later point in time for which binary log data is available (and accept possible data loss on the target).
- Resynchronize the source and target tables, and then start the Oracle GoldenGate environment over again.

To determine where the Extract checkpoints are, use the INFO EXTRACT command. For more information, see INFO EXTRACT in *Command Line Interface Reference for Oracle GoldenGate*.

Setting Logging Parameters

To capture from the MySQL transaction logs, the Oracle GoldenGate Extract process must be able to find the index file, which contains the paths of all binary log files.

Extract expects that all of the table columns are in the binary log. As a result, only binlog_row_image set as full is supported and this is the default. Other values of binlog row image are not supported.



Oracle recommends that the binary log is retained for at least 24 hours.

In MySQL 5.7, the <code>server_id</code> option must be specified along with <code>log-bin</code>, otherwise the server will not start. For MySQL 8.0, the <code>server_id</code> is enabled by default.

Extract checks the following parameter settings to get this index file path:

1.

2. Extract TRANLOGOPTIONS parameter with the ALTLOGDEST option. If this parameter specifies a location for the log index file, Extract accepts this location over any default that is specified in the MySQL Server configuration file. When ALTLOGDEST is used, the binary log index file must also be stored in the specified directory. This parameter should be used if the MySQL configuration file does not specify the full index file path name, specifies an incorrect location, or if there are multiple installations of MySQL on the same machine. From Oracle GoldenGate 21c onward, ALTLOGDEST parameter is optional. When ALTLOGDEST is not specified, the binary log index and binary log filepaths will be fetched from the database directly.Please note: The paths thus fetched are also subject to same accessibilitychecks as in the existing process.



To specify the index file path with TRANLOGOPTIONS with ALTLOGDEST, use the following command format on Windows:

TRANLOGOPTIONS ALTLOGDEST "C:\Program Files\MySQL\logs\binlog.index"

On Linux, use this format:

TRANLOGOPTIONS ALTLOGDEST "/mnt/rdbms/mysql/data/logs/binlog.index"

To capture from a remote server or in case of remote capture, you only need to specify the REMOTE option instead of the index file path on the remote server. For remote capture on both Windows and Linux, specify the following in the Extract parameter file:

TRANLOGOPTIONS ALTLOGDEST REMOTE

- 3. The MySQL Server configuration file: The configuration file stores default startup options for the MySQL server and clients. On Windows, the name of the configuration file is my.ini. On other platforms, it is my.cnf. In the absence of TRANLOGOPTIONS with ALTLOGDEST, Extract gets information about the location of the log files from the configuration file. However, even with ALTLOGDEST, these Extract parameters must be set correctly:
 - binlog-ignore-db=oggddl: This prevents DDL logging history table entries in the binlog and is set in the my.cnf or my.ini file.
 - log-bin: This parameter is used to enable binary logging. This parameter also specifies the location of the binary log index file and is a required parameter for Oracle GoldenGate, even if ALTLOGDEST is used. If log-bin is not specified, binary logging will be disabled and Extract returns an error.
 - log-bin-index: This parameter specifies the location of the binary log index. If it is not used, Extract assumes that the index file is in the same location as the log files. If this parameter is used and specifies a different directory from the one that contains the binary logs, the binary logs must not be moved once Extract is started.
 - max binlog size: This parameter specifies the size, in bytes, of the binary log file.

Note:

The server creates a new binary log file automatically when the size of the current log reaches the <code>max_binlog_size</code> value, unless it must finish recording a transaction before rolling over to a new file.

binlog_format: This parameter sets the format of the logs. It must be set to the value
of ROW, which directs the database to log DML statements in binary format. From
Oracle GoldenGate 19c onward, Extract silently ignores the binlog events that are not
written in the ROW format instead of abending when it detects a binlog_format other
than ROW.



Note:

MySQL binary logging does not allow logging to be enabled or disabled for specific tables. It applies globally to all tables in the database.

mysql.rds_set_configuration: When capturing from MySQL Amazon RDS instance, you need to call the mysql.rds_set_configuration stored procedure on MySQL command line, to retain the binary logs for a specific duration. By default, the default value of binlog_retention_hours for MySQL Amazon RDS is set to NULL, which implies that the binary logs are not retained.

The following example shows the command to preserve the binary log for 24 hours:

```
mysql > call mysql.rds set configuration('binlog retention hours', 24);
```

To locate the configuration file, Extract checks the MYSQL_HOME environment variable: If MYSQL_HOME is set, Extract uses the configuration file in the specified directory. If MYSQL_HOME is not set, Extract queries the information_schema.global_variables table to determine the MySQL installation directory. If a configuration file exists in that directory, Extract uses it.

4. For MariaDB version 10.2 and later, Oracle GoldenGate works in the same way as for MySQL but a new variable needs to be configured in the my.cnf or my.ini file. The variable that needs to be added is "binlog-annotate-row-events=OFF". Restart MariaDB after configuring this variable and then start the Extract process.

Adding Host Names

Oracle GoldenGate gets the name of the database it is supposed to connect to from the SOURCEDB parameter. A successful connection depends on the localhost entry being properly configured in the system host file. To avoid issues that arise from improper local host configuration, you can use SOURCEDB in the following format:

```
SOURCEDB dbname@hostname:port, USERID mysqluser, PASSWORD welcome
```

The dbname is the name of the MySQL instance, hostname is the name or IP address, port is the port number of the MySQL instance. If using an unqualified host name, that name must be properly configured in the DNS database. Otherwise, use the fully qualified host name, for example myhost.company.com.

Setting the Session Character Set

The GGSCI, Extract and Replicat processes use a session character set when connecting to the database. For MySQL, the session character set is taken from the SESSIONCHARSET option of SOURCEDB and TARGETDB. Make certain you specify a session character set in one of these ways when you configure Oracle GoldenGate.

Preparing Tables for Processing

This section describes how to prepare the tables for processing. Table preparation requires these tasks:

Assigning Row Identifiers



- Limiting Row Changes in Tables That Do Not Have a Key
- Triggers and Cascade Constraints Considerations

Assigning Row Identifiers

Oracle GoldenGate requires some form of unique row identifier on the source and target tables to locate the correct target rows for replicated updates and deletes.

- How Oracle GoldenGate Determines the Kind of Row Identifier to Use
- Tables with a Primary Key Derived from a Unique Index
- How to Specify Your Own Key for Oracle GoldenGate to Use

How Oracle GoldenGate Determines the Kind of Row Identifier to Use

Unless a KEYCOLS clause is used in the TABLE or MAP statement, Oracle GoldenGate selects a row identifier to use in the following order of priority:

- Primary key
- 2. First unique key alphanumerically that does not contain a timestamp or non-materialized computed column.
- 3. If none of the preceding key types exist (even though there might be other types of keys defined on the table) Oracle GoldenGate constructs a pseudo key of all columns that the database allows to be used in a unique key, excluding those that are not supported by Oracle GoldenGate in a key or those that are excluded from the Oracle GoldenGate configuration.



If there are other, non-usable keys on a table or if there are no keys at all on the table, Oracle GoldenGate logs an appropriate message to the report file. Constructing a key from all of the columns impedes the performance of Oracle GoldenGate on the source system. On the target, this key causes Replicat to use a larger, less efficient WHERE clause.

Tables with a Primary Key Derived from a Unique Index

In the absence of a primary key on a table, MySQL will promote a unique index to primary key if the indexed column is \mathtt{NOT} \mathtt{NULL} . If there are more than one of these not-null indexes, the first one that was created becomes the primary key. To avoid Replicat errors, create these indexes in the same order on the source and target tables.

For example, assume that source and target tables named ggvam.emp each have columns named first, middle, and last, and all are defined as NOT NULL. If you create unique indexes in the following order, Oracle GoldenGate will abend on the target because the table definitions do not match.

Source:

```
mysql> create unique index uq1 on ggvam.emp(first);
mysql> create unique index uq2 on ggvam.emp(middle);
mysql> create unique index uq3 on ggvam.emp(last);
```



Target:

```
mysql> create unique index uq1 on ggvam.emp(last);
mysql> create unique index uq2 on ggvam.emp(first);
mysql> create unique index uq3 on ggvam.emp(middle);
```

The result of this sequence is that MySQL promotes the index on the source "first" column to primary key, and it promotes the index on the target "last" column to primary key. Oracle GoldenGate will select the primary keys as identifiers when it builds its metadata record, and the metadata will not match. To avoid this error, decide which column you want to promote to primary key, and create that index first on the source and target.

How to Specify Your Own Key for Oracle GoldenGate to Use

If a table does not have one of the preceding types of row identifiers, or if you prefer those identifiers not to be used, you can define a substitute key if the table has columns that always contain unique values. You define this substitute key by including a KEYCOLS clause within the Extract TABLE parameter and the Replicat MAP parameter. The specified key will override any existing primary or unique key that Oracle GoldenGate finds.

Limiting Row Changes in Tables That Do Not Have a Key

If a target table does not have a primary key or a unique key, duplicate rows can exist. In this case, Oracle GoldenGate could update or delete too many target rows, causing the source and target data to go out of synchronization without error messages to alert you. To limit the number of rows that are updated, use the DBOPTIONS parameter with the LIMITROWS option in the Replicat parameter file. LIMITROWS can increase the performance of Oracle GoldenGate on the target system because only one row is processed.

Triggers and Cascade Constraints Considerations

Triggers

Disable triggers on the target tables, or alter them to ignore changes made by the Oracle GoldenGate database user. Oracle GoldenGate replicates DML that results from a trigger. If the same trigger gets activated on the target table, then it becomes redundant because of the replicated version, and the database returns an error.

Cascade Constraints Considerations

Cascading updates and deletes captured by Oracle GoldenGate are not logged in binary log, so they are not captured. This is valid for both MySQL and MariaDB. For example, when you run the delete statement in the parent table with a parent child relationship between tables, the cascading deletes (if there are any) happens for child table, but they are not logged in binary log. Only the delete or update record for the parent table is logged in the binary log and captured by Oracle GoldenGate.

See https://mariadb.com/kb/en/replication-and-foreign-keys/ and https://dev.mysql.com/doc/refman/8.0/en/innodb-and-mysql-replication.html for details.

To properly handle replication of cascading operations, it is recommended to disable cascade deletes and updates on the source and code your application to explicitly delete or update the child records prior to modifying the parent record. Alternatively, you must ensure that the target parent table has the same cascade constraints configured as the source parent table, but this could lead to an out-of-sync condition between source and target, especially in cases of bidirectional replication.



Changing the Log-Bin Location

Modifying the binary log location by using the log-bin variable in the MySQL configuration file might result in two different path entries inside the index file, which could result in errors. To avoid any potential errors, change the log-bin location by doing the following:

- 1. Stop any new DML operations.
- Let the extract finish processing all of the existing binary logs. You can verify this by noting when the checkpoint position reaches the offset of the last log.
- 3. After Extract finishes processing the data, stop the Extract group and, if necessary, back up the binary logs.
- 4. Stop the MySQL database.
- 5. Modify the log-bin path for the new location.
- 6. Start the MySQL database.
- To clean the old log name entries from index file, use flush master or reset master (based on your MySQL version).
- 8. Start Extract.

Configuring Bi-Directional Replication

In a bi-directional configuration, there are Extract and Replicat processes on both the source and target systems to support the replication of transactional changes on each system to the other system. To support this configuration, each Extract must be able to filter the transactions applied by the local Replicat, so that they are not recaptured and sent back to their source in a continuous loop. Additionally, AUTO_INCREMENT columns must be set so that there is no conflict between the values on each system.

- Configure Oracle GoldenGate for high availability or active-active replication according to the instructions in the Propagating DDL in Active-Active (Bidirectional) Configurations in Using Oracle GoldenGate for Oracle Database.
- 2. To filter out Replicat operations in a bi-directional configuration so that the applied operations are not captured and looped back to the source again, take the following steps on each MySQL database:
 - Configure each Replicat process to use a checkpoint table. Replicat writes a
 checkpoint to this table at the end of each transaction. You can use one global
 checkpoint table or one per Replicat process See Overview of Replicat in
 Understanding Oracle GoldenGate.
 - Specify the name of the checkpoint table with the FILTERTABLE option of the TRANLOGOPTIONS parameter in the Extract parameter file. The Extract process will ignore transactions that end with an operation to the specified table, which should only be those of Replicat.



Although optional for other supported databases as a means of enhancing recovery, the use of a checkpoint table is required for MySQL when using bidirectional replication (and likewise, will enhance recovery).



3. Edit the MySQL server configuration file to set the auto_increment_increment and auto_increment_offset parameters to avoid discrepancies that could be caused by the bidirectional operations. The following illustrates these parameters, assuming two servers: ServerA and ServerB.

ServerA:

```
auto-increment-increment = 2
auto-increment-offset = 1

ServerB:
auto-increment-increment = 2
auto-increment-offset = 2
```

Configuring MySQL for Remote Capture

Oracle GoldenGate remote capture for MySQL, Amazon RDS for MySQL, Amazon Aurora MySQL, Azure Database for MySQL are used to capture transaction log data from a database located remotely to the Oracle GoldenGate installation.

Database Server Configuration

For remote capture to work, configure the MySQL server as follows:

1. Grant access permissions to the Oracle GoldenGate remote capture user.

Run the following statements against the remote database to create the user and grant the permissions needed for remote capture.

```
mysql > CREATE USER 'username'@'host' IDENTIFIED BY 'Password';
mysql > GRANT ALL PRIVILEGES ON *.* TO 'username'@'host' WITH GRANT
OPTION;
mysql > FLUSH PRIVILEGES;
```

2. The server_id value of the remote MySQL server should be greater than 0. This value can be verified by issuing the following statement on the MySQL remote server:

```
mysql > show variables like 'server id';
```

If the $server_id$ value is 0, modify the my.cnf configuration file to set to a value greater than 0.

Oracle GoldenGate Configuration

Oracle GoldenGate configuration has the following steps:

Provide the remote database's connection information in the Extract's parameter file.

```
SOURCEDB remotedb@mysqlserver.company.com:port, USERID username, PASSWORD password
```

Add the following parameter to the Extract's parameter file, after the connection information.

```
TRANLOGOPTIONS ALTLOGDEST REMOTE
```



Limitations of Oracle GoldenGate Remote Capture for MySQL

Co-existence of Oracle GoldenGate for MySQL remote capture with the MySQL's native replication slave is supported with following conditions and limitations:

 The native replication slave should be assigned a different server_id than the currently running slaves. The slave server_id values can be seen using the following MySQL command on the master server.

```
mysql> show slave hosts;
```

- If the Oracle GoldenGate capture abends with error "A slave with the same server_uuid or server_id as this slave has connected to the master", then change the capture's name and restart the capture.
- If the native replication slave dies with the error "A slave with the same server_uuid or server_id as this slave has connected to the master", then change the native replication slave's server id and restart it.
- Remote capture is supported only on the Linux 64-bit platform and not on Windows. But Oracle GoldenGate remote capture on Linux can capture from a MySQL database running on remote Windows machine.

Configuring a Two-way SSL Connection in MySQL Capture and Delivery

To use the two way SSL in Oracle GoldenGate MySQL capture and delivery, you need to supply the full paths of the certificate authority (ca.pem), the client certificate (client-cert.pem) and the client key (client-key.pem) files to the capture and delivery.

To know more about generating the certificate files, see:

https://dev.mysql.com/doc/refman/5.7/en/creating-ssl-rsa-files-using-mysql.html

You need to provide these paths in the Extract and Replicat parameter files using the SETENV parameter.

Following are the SETENV environment parameters to set the two-way SSL connection:

- OGG MYSQL OPT SSL CA: Sets the full path of the certification authority.
- OGG MYSQL OPT SSL CERT: Sets the full path of the client certificate.
- OGG MYSQL OPT SSL KEY: Sets the full path of the client key.

In the following example, the MySQL SSL certificate authority, client certificate, and client key paths are set to the Oracle GoldenGate MySQL Extract and Replicat parameter:

```
SETENV (OGG_MYSQL_OPT_SSL_CA='/var/lib/mysql.pem')
SETENV (OGG_MYSQL_OPT_SSL_CERT='/var/lib/mysql/client-cert.pem')
SETENV (OGG MYSQL OPT SSL KEY='/var/lib/mysql/client-key.pem')
```



Capturing using a MySQL Replication Slave

You can configure a MySQL replication slave to capture the master's binary log events from the slave.

Typically, the transactions applied by the slave are logged into the relay logs and not into the slave's <code>binlog</code>. For the slave to write transactions in its <code>binlog</code>, that it receives from the master , you must start the replication slave with the <code>log-slave-updates</code> option as 1 in <code>my.cnf</code> in conjunction with the other binary logging parameters for Oracle GoldenGate. After the master's transactions are in the slave's <code>binlog</code>, you can set up a regular Oracle GoldenGate capture on the slave to capture and process the slave's <code>binlog</code>.

Other Oracle GoldenGate Parameters for MySQL

The following parameters may be of use in MySQL installations, and might be required if non-default settings are used for the MySQL database. Other Oracle GoldenGate parameters will be required in addition to these, depending on your intended business use and configuration.

Parameter	Description
DBOPTIONS with CONNECTIONPORT port_number	Required to specify to the VAM the TCP/IP connection port number of the MySQL instance to which an Oracle GoldenGate process must connect if MySQL is not running on the default of 3306. DBOPTIONS CONNECTIONPORT 3307
DBOPTIONS with HOST host_id	Specifies the DNS name or IP address of the system hosting MySQL to which Replicat must connect.
DBOPTIONS with ALLOWLOBDATATRUNCATE	Prevents Replicat from abending when replicated LOB data is too large for a target MySQL CHAR, VARCHAR, BINARY or VARBINARY column.
SOURCEDB with USERID and PASSWORD	Specifies database connection information consisting of the database, user name and password to use by an Oracle GoldenGate process that connects to a MySQL database. If MySQL is not running on the default port of 3306, you must specify a complete connection string that includes the port number: SOURCEDB dbname@hostname:port, USERID user, PASSWORD password.Example:
	SOURCEDB mydb@mymachine:3307, USERID myuser, PASSWORD mypassword
	If you are not running the MySQL database on port 3306, you must also specify the connection port of the MySQL database in the DBLOGIN command when issuing commands that affect the database through GGSCI:
	DBLOGIN SOURCEDB dbname@hostname:port, USERID user, PASSWORD password
	For example:
	GGSCI> DBLOGIN SOURCEDB mydb@mymachine:3307, USERID myuser, PASSWORD mypassword



Parameter	Description
SQLEXEC	To enable Replicat to bypass the MySQL connection timeout, configure the following command in a SQLEXEC statement in the Replicat parameter file.
	SQLEXEC "select CURRENT_TIME();" EVERY n MINUTES
	Where : n is the maximum interval after which you want Replicat to reconnect. The recommended connection timeout 31536000 seconds (365 days).



Parameter Description For heartbeattable to work in MySQL 5.7, MySQL global variable Global variable sql mode sql_mode should not have NO_ZERO_IN_DATE, NO_ZERO_DATE. In the following example sql mode includes NO ZERO IN DATE, NO ZERO DATE values: mysql> show variables like '%sql mode%';+------| Variable_name | Value +-----| sql mode ONLY FULL GROUP BY, STRICT TRANS TABLES, NO ZERO IN DA TE, NO ZERO DATE, ERROR FOR DIVISION BY ZERO, NO AUTO CREATE USER, NO ENGINE SUBST ITUTION | 1 row in set (0.00 sec) These values must be removed by issuing the following command: mysql> Set global sql mode='ONLY FULL GROUP BY, STRICT TRANS TABLES, ERR OR FOR DIVISION BY ZERO, NO AUTO CREATE USER, NO ENGINE SUBSTITUTION'; Query OK, 0 rows affected, 1 warning (0.00 sec) mysql> show variables like '%sql mode%'; +----| Variable name | Value +-----_____+ | sql mode ONLY FULL GROUP BY, STRICT TRANS TABLES, ERROR FOR DIV ISION BY ZERO, NO AUTO CREA

Parameter	Description	
	TE_USER,NO_ENGINE_SUBSTITUTION +	
	1 row in set (0.01 sec)	

Positioning Extract to a Specific Start Point

You can position the Extract to a specific start point in the transaction logs using the ADD/ALTER EXTRACT commands:

{ADD | ALTER EXTRACT} group, LOGNUM log num, LOGPOS log pos

- *group* is the name of the Oracle GoldenGate Extract group for which the start position is required.
- LOGNUM is the log file number. For example, if the required log file name is test.000034, the LOGNUM value is 34. Extract will search for this log file.
- LOGPOS is an event offset value within the log file that identifies a specific transaction record. Event offset values are stored in the header section of a log record. To position at the beginning of a binlog file, set the LOGPOS as 0.

In MySQL logs, an event offset value can be unique only within a given binary file. The combination of the position value and a log number will uniquely identify a transaction record. Maximum Log number length is 8 bytes unsigned integer and Maximum Log offset length is 8 bytes unsigned integer. Log number and Log offset are separated by a pipe ('|') delimiter. Transactional records available after this position within the specified log will be captured by Extract. In addition, you can position an Extract using a timestamp.



Using DDL Replication

Learn how to install, use, configure, and remove DDL replication. **Data Definition Language (DDL)** statements (operations) are used to define MySQL database structures or schema. You can use these DDL statements for data replication between MySQL source and target databases. MySQL DDL specifics are found in the MySQL documentation at https://dev.mysql.com/doc/.

Topics:

- Plug-in Based DDL Configuration Prerequisites and Considerations
- Installing DDL Replication
- Using the Metadata Server
- Using DDL Filtering for Replication
- Troubleshooting Plug-in Based DDL Replication
- Uninstalling Plug-In Based DDL Replication

Plug-in Based DDL Configuration Prerequisites and Considerations

This is an older approach to performing DDL Replication. The prerequisites for configuring DDL replication are as follows:

- DDL replication is supported for MySQL 5.7.
- Remote capture for MySQL 5.7 doesn't support DDL replication.
- Oracle GoldenGate DDL replication uses two plug-ins as a shared library, ddl_rewriter and ddl_metadata, which must be installed on your MySQL server before Oracle GoldenGate replication starts.
- The standalone application, Oracle GoldenGate metadata_server, must be running to capture the DDL metadata.
- The history table under the new oggddl database (oggddl.history). This metadata history table is used to store and retrieve the DDL metadata history. The history table records must be ignored from being logged into the binary log so you must specify binlog-ignore-db=oggddl in the my.cnf file.
- You should not manually drop the oggddl database or the history table because all DDL statements that run after this event will be lost.
- You should not stop the metadata_server during DDL capture as all the DDL statements
 that run after this event will be lost.
- You should not manually remove the ddl_rewriter and the ddl_metadata plugins during DDL capture because all DDL statements that run after this event will be lost.

DDL executed within the stored procedure is not supported. For example, a DDL executed
as in the following is not supported.

```
CREATE PROCEDURE atssrc.generate data()
BEGIN
DECLARE i INT DEFAULT 0;
WHILE i < 800 DO
SET i = i + 1;
IF (i = 100) then
alter table atssrc. `ddl6` add col2 DATE after id;
ELSEIF (i = 200) then
alter table atssrc. `ddl6` add col3 DATETIME after datetime;
ELSEIF (i = 300) then
alter table atssrc.`ddl6` add `col4` timestamp NULL DEFAULT NULL after
channel;
ELSEIF (i = 400) then
alter table atssrc. `ddl6` add col5 YEAR after value;
END IF;
END WHILE;
END$$
DELIMITER ;
call atssrc.generate data();
```

 By design, the heartbeat table DDLs are ignored by the capture and you should create the heartbeat tables manually at the target.

Installing DDL Replication

To install DDL replication, you run the installation script that is provided with Oracle GoldenGate as the replication user. This user must have Create, Insert,Select, Delete, Drop, and Truncate database privileges. Additionally, this user must have write permission to copy the Oracle GoldenGate plugin in the MySQL plugin directory. For example, the MySQL plugin are typically in /usr/lib64/mysql/plugin/.

The installation script options are install, uninstall, start, stop, and restart.

The command to install DDL replication uses the install option, user id, password, and port number respectively:

```
bash-3.2$ ./ddl install.sh install-option user-id password port-number
```

For example:

```
bash-3.2$ ./ddl install.sh install root welcome 3306
```

The DDL replication installation script completes the following tasks:

- Ensures that you have a supported MySQL server version installed. DDL replication is supported for MySQL 5.7.10 and greater.
- Locates the MySQL plugin directory.
- 3. Ensures that the ddl_rewriter, ddl_metadata plugins and the metadata_server files exist. If these files are not found, then an error message appears and the installation exits.

- Ensures that the plugins are already installed. If installed, the script exits with a message requesting you to uninstall first and then reinstall.
- Stops the metadata_server if it is running.
- 6. Deletes the oggddl.history table if it exists.
- 7. Starts the metadata server as a daemon process.
- 8. Installs the ddl rewriter and ddl metadata plugins.

Using the Metadata Server

You can use the following options with the metadata server:

- You must have the Oracle GoldenGate metadata_server running to capture the DDL metadata.
- Run the install script with start option to start the metadata server.
- Run the install script with stop option to stop the metadata server.
- Run the install script with restart option to stop the running metadata server and start again.
- Oracle GoldenGate DDL replication uses two plugins as a shared library, ddl_rewriter and ddl_metadata, both of which must be installed on your MySQL server before Oracle GoldenGate replication starts.
- The oggddl.history metadata history table is used to store and retrieve the DDL metadata history.

There is a single history table and metadata server for each MySQL server. If you want to issue and capture DDLs from multiple instances of an Extract process on the same database server at the same time, there is a possibility of conflict between accessing and populating the metadata history table. Oracle recommends that you do not run and capture DDLs using multiple Extract instances on the same MySQL server.

Using DDL Filtering for Replication

The following options are supported for MySQL DDL replication:

Option	Description
DDL INCLUDE OPTYPE CREATE OBJTYPE TABLE;	Include create table.
DDL INCLUDE OBJNAME ggvam.*	Include tables under the ggvamdatabase.
DDL EXCLUDE OBJNAME ggvam.emp*;	Exclude all the tables under the ggvam database and table name starting with the empwildcard.
DDL INCLUDE INSTR 'XYZ'	Include DDL that contains this string.
DDL EXCLUDE INSTR 'WHY'	Excludes DDL that contains this string.
DDL INCLUDE MAPPED	MySQL DDL uses this option and should be used as the default for Oracle GoldenGate MySQL DDL replication. DDL INCLUDE ALL and DDL are not supported.
DDL EXCLUDE ALL	Default option.



For a full list of options, see DDL in Reference for Oracle GoldenGate.

Using DDL Statements and Options

- INCLUDE (default) means include all objects that fit the rest of the description. EXCLUDE
 means to omit items that fit the description. Exclude rules take precedence over include
 rules.
- OPTYPE specifies the types of operations to be included or excluded. You can use CREATE and ALTER. Multiple OPTYPE can be specified using parentheses. For example, optype (create, alter). The asterisk (*) wildcard can be specified to indicate all operation types, and this is the default.
- OBJTYPE specifies the TABLE operations to include or exclude. The wildcard can be specified to indicate all object types, and this is the default.
- OBJNAME specifies the actual object names to include or exclude. For example, eric.*.
 Wildcards are specified as in other cases where multiple tables are specified. The default is *.
- String indicates that the rule is true if any of the strings in stringspec are present (or false if excludestring is specified and the stringspec is present). If multiple string entries are made, at least one entry in each stringspec must be present to make the rule evaluate true.

For example:

```
ddlops string ("a", "b"), string ("c") evaluates true if string "a" OR "b" is present, AND string "c" is present
```

- local is specified if you want the rule to apply only to the current Extract trail (the Extract trail to which the rule applies must precede this ddlops specification).
- The semicolon is required to terminate the parameter entry.

For example:

```
ddl optype (create, drop), objname (eric.*);
ddl exclude objname (eric.tab*);
exttrail a;
exttrail b;
ddl optype (create), objname (joe.*), string ("abc", "xyz") local;
ddl optype (alter), objtype (index);
```

In this preceding example, the <code>exttrail</code> a gets creates and drops for all objects that belong to <code>eric</code>, except for objects that start with <code>tab</code>, <code>exttrail</code> a also gets all alter index statements, unless the index name begins with <code>tab</code> (the rule is global even though it's included in <code>exttrail</code> b). <code>exttrail</code> b gets the same objects as <code>a</code>, and it also gets all creates for objects that belong to <code>joe</code> when the string <code>abcor</code> <code>xyz</code> is present in the DDL text. The <code>ddlops.c</code> module stores all DDL operation parameters and executes related rules.

Additionally, you can use the DDLOPTIONS parameter to configure aspects of DDL processing other than filtering and string substitution. You can use multiple DDLOPTIONS statements and Oracle recommends using one. If you are using multiple DDLOPTIONS statements, then make each of them unique so that one does not override the other. Multiple DDLOPTIONS statements are executed in the order listed in the parameter file.

See DDL and DDLOPTIONS.

Troubleshooting Plug-in Based DDL Replication

Plug-in based DDL replication relies on a metadata history table and the metadata plugin and server. To troubleshoot when DDL replication is enabled, the history table contents and the metadata plugin server logs are required.

You can use the <code>mysqldump</code> command to generate the history table dump using one of the following examples:

```
mysqldump [options] database [tables]
mysqldump [options] --databases [options] DB1 [DB2 DB3...]
mysqldump [options] --all-databases [options]
```

For example, bash-3.2\$ mysqldump -uroot -pwelcome oggddl history > outfile

The metadata plugins and server logs are located in the MySQL and Oracle GoldenGate installation directories respectively.

If you find an error in the log files, you need to ensure that the metadata server is running.

Uninstalling Plug-In Based DDL Replication

If you no longer want to capture the DDL events, then you can use the same install script and select the uninstall option to disable the DDL setup. Also, any Extract with DDL parameters should be removed or disabled. If you want to capture the DDL again, you can run the install script again. You should take care when multiple instances of the capture process is running on the same instance of your MySQL server. The DDL setup should *not* be disturbed or uninstalled when multiple capture processes are running and when at most one capture is designed to capture the DDL statement.

Use the installation script with the uninstall option to uninstall DDL Replication. For example:

```
bash-3.2$ ./ddl install.sh uninstall root welcome 3306
```

The script performs the following tasks:

- Uninstalls the ddl rewriter and ddl metadata plugins.
- 2. Deletes the oggddl.history table if exists.
- 3. Removes the plugins from MySQL plugin directory.
- Stops the metadata server if it is running.



Part VI

Using Oracle GoldenGate for PostgreSQL

Oracle GoldenGate for PostgreSQL supports capture and delivery of initial load and transactional data for supported PostgreSQL database versions.

Oracle GoldenGate for PostgreSQL supports the mapping, filtering, and transformation of source data, unless noted otherwise in this document, as well as replicating data derived from other source databases supported by Oracle GoldenGate, into PostgreSQL databases.

This part describes the tasks for configuring and running Oracle GoldenGate for PostgreSQL.

Topics:

- Understanding What's Supported for PostgreSQL
- · Preparing the Database for Oracle GoldenGate
- Configuring Extract
- · Configuring Replicat
- Additional Considerations



15

Understanding What's Supported for PostgreSQL

This chapter contains information on supported features for Oracle GoldenGate for PostgreSQL:

Topics:

- Supported Databases
- Details of Supported PostgreSQL Data Types
- Supported Objects and Operations for ProtgreSQL

Supported Databases

The following are supported databases and limitations for Oracle GoldenGate for PostgreSQL:

- Only user databases are supported for capture and delivery.
- Oracle GoldenGate does not support capture from archived logs.
- Capture and delivery are not supported against replica, standby databases.
- High Availability:
 - Oracle GoldenGate Extract does *not* support seamless role transitioning from a
 primary to a secondary Extract with PostgreSQL high availability configurations.
 However, manual procedural operations could be followed to provide continuity from
 the new primary Extract.
 - For more information, see the details available in the my Oracle Support note, Oracle GoldenGate Procedures for PostgreSQL HA Failover (Doc ID 2818379.1).

Details of Supported PostgreSQL Data Types

This topic describes data types that are supported by Oracle GoldenGate, as well as those that are not supported.

- Supported PostgreSQL Data Types
- Non-Supported PostgreSQL Data Types

Supported PostgreSQL Data Types

Here's a list of PostgreSQL data types that Oracle GoldenGate supports along with the limitations of this support.

- bigint
- bigserial
- bit(n)
- bit varying(n)

- boolean
- bytea
- char (n)
- cidr
- date
- decimal
- double precision
- inet
- integer
- interval
- json
- jsonb
- macaddr
- macaddr8
- money
- numeric
- real
- serial
- smallint
- smallserial
- text
- time with/without timezone
- timestamp with/without timezone
- uuid
- varchar(n)
- varbit
- xml

Limitations of Support

- If columns of char, varchar, text, or bytea data types are part of a primary or unique key, then the maximum individual lengths for these columns must not exceed 8191 bytes.
- real, double, numeric, decimal: NaN input values are not supported.
- The following limitations apply to bit/varbit data types:
 - They are supported up to 4k in length. For lengths greater than 4k the data is truncated and only the lower 4k bits are captured.
 - The source bit(n) column can be applied only onto a character type column on a non-PostgreSQL target and can be applied onto a char type or a bit/varbit column on PostgreSQL target.



- The following limitations are applicable to both timestamp with time zone and timestamp without time zone:
 - The timestamp data with BC or AD tags in the data is not supported.
 - The timestamp data older than 1883-11-18 12:00:00 is not supported.
 - The timestamp data with more than 4 digits in the YEAR component is not supported.
 - Infinity/-Infinity input strings for timestamp columns are not supported.
- The following are the limitations when using interval:
 - The capture of mixed sign interval data from interval type columns is not supported. You can use DBOPTIONS ALLOWNONSTANDARDINTERVALDATA in the Extract parameter file to capture the mixed sign interval data (or any other format of interval data, which is not supported by Oracle GoldenGate) as a string (not as standard interval data).

The following are a few examples of data that gets written to the trail file, on using the DBOPTIONS ALLOWNONSTANDARDINTERVALDATA in the Extract param file:

- +1026-9 +0 +0:0:22.000000 is interpreted as 1026 years, 9 months, 0 days, 0 hours,
 0 minutes, 22 seconds.
- -0-0 -0 -8 is interpreted as 0 years, 0 months, 0 days, -8 hours.
- +1-3 +0 +3:20 is interpreted as 1 year, 3 months, 0 days, 3 hours, 20 minutes.
- Replicat: If the source interval data was captured using DBOPTIONS
 ALLOWNONSTANDARDINTERVALDATA and written as a string to the trail, the corresponding source column is allowed to be mapped to either a char or a binary type column on the target.
- date limitations are:
 - The date data with BC or AD tags in the data is not supported.
 - Infinity/-Infinity input strings for date columns are not supported.
- Columns of text, json, xml, bytea, char (>8191), varchar (>8191) are treated as LOB columns and have the following limitations:
 - When using GETUPDATEBEFORES, the before image of LOB columns is never logged.
 - When using NOCOMPRESSUPDATES, LOB columns are logged in the after image only if they were modified.
- The support of range and precision for floating-point numbers depends on the host machine. In general, the precision is accurate to 16 significant digits, but you should review the database documentation to determine the expected approximations. Oracle GoldenGate rounds or truncates values that exceed the supported precision.

Non-Supported PostgreSQL Data Types

Oracle GoldenGate for PostgreSQL does not support the following data types:

- Arrays
- box
- circle
- Composite Types



- Domain Types
- Enumerated Types
- line
- lsea
- Object Identifiers Types
- path
- pg lsn
- · pg snapshot
- point
- polygon
- Pseudo-Types
- Range Types
- tsquery
- tsvector
- User-defined Types (UDTs)
- Extensions and Additional Supplied Modules listed at: https://www.postgresql.org/docs/ current/contrib.html are not supported by Oracle GoldenGate unless explicitly listed under Supported PostgreSQL Data Types.

Note:

If the Extract parameter file contains a table with unsupported data types, the Extract will stop with an error message. To resume replication, remove the table from the Extract file or remove the column from the table with an unsupported data type.

Note:

If an Extension or Additional Supplied Module is supported, it will be explicitly added to the Supported PostgreSQL data types list.

Supported Objects and Operations for ProtgreSQL

This topic describes objects and operations that are supported by Oracle GoldenGate.

- Oracle GoldenGate for PostgreSQL only supports DML operations (Insert/Update/Deletes).
 DDL replication is not supported.
- Oracle GoldenGate for PostgreSQL supports replication of truncate operations beginning with PostgreSQL 11 and above, and requires the GETTRUNCATES parameter in Extract and Replicat.
- Case Sensitive/Insensitive names usage includes:



- Unquoted names are case-insensitive and are implicitly lowercase. For example,
 CREATE TABLE MixedCaseTable and SELECT * FROM mixedcasetable are equivalent.
- Quoted table and column names are case-sensitive and need to be listed correctly in Extracts and Replicats and with Oracle GoldenGate commands. For example, TABLE appschema."MixedCaseTable" and ADD TRANDATA appschema."MixedCaseTable" would be required to support a case-sensitive table name.
- Tables and Views
- Sequences and Identity Columns

Tables and Views

Tables to be included for capture and delivery must meet the following requirements and must only include data types listed under *Supported PostgreSQL Data Types*.

- Oracle GoldenGate for PostgreSQL supports capture of transactional DML from user tables, and delivery to user tables.
- · Oracle GoldenGate for PostgreSQL supports delivery to partitioned tables.
- Globalization is supported for object names (table /schema/database/column names) and column data.

Limitations:

- Oracle GoldenGate for PostgreSQL does not support capture and delivery for views.
- Oracle GoldenGate for PostgreSQL does not support capture from partitioned tables.

Sequences and Identity Columns

- Sequences are supported on source and target tables for unidirectional, bidirectional, and multi- directional implementations.
- Identity columns created using the GENERATED BY DEFAULT AS IDENTITY clause are supported on source and target tables, for unidirectional, bidirectional, and multi-directional implementations.
- Identity columns created using the GENERATED ALWAYS AS IDENTITY clause are not supported in target database tables and the Identity property should be removed from target tables or changed to GENERATED BY DEFAULT AS IDENTITY.
- For bidirectional and multi-directional implementations, define the Identity columns and sequences with an INCREMENT BY value equal to the number of servers in the configuration, with a different MINVALUE for each one.

For example, MINVALUE /INCREMENT BY values for a bidirectional, two-database configuration would be as follows:

Database1, set the MINVALUE at 1 with an INCREMENT BY of 2.

Database2, set the MINVALUE at 2 with an INCREMENT BY of 2.

For example, MINVALUE /INCREMENT BY values for a multi-directional, three-database configuration would be as follows:

Database1, set the MINVALUE at 1 with an INCREMENT BY of 3.

Database2, set the MINVALUE at 2 with an INCREMENT BY of 3.

Database3, set the MINVALUE at 3 with an INCREMENT BY of 3.



16

Preparing the Database for Oracle GoldenGate

This chapter describes how to prepare your PostgreSQL database and environment for Oracle GoldenGate.

Topics:

- Database Configuration
- Establishing Oracle GoldenGate Credentials
- Configuring a Database Connection
- Preparing Tables for Processing

Database Configuration

To support Oracle GoldenGate, the following parameters in the PostgreSQL database configuration file, \$PGDATA/postgresql.conf, needs to be configured.

 For remote connectivity of an Extract or Replicat, set the PostgreSQL listen_addresses to allow for remote database connectivity. For example:

```
listen addresses=remotehost ip address
```



Ensure that client authentication is set to allow connections from an Oracle GoldenGate host by configuring the pg_hba.conf file. For more information, refer to this document: The pg_hba.conf File

• To support Oracle GoldenGate capture, write-ahead logging must be set to logical, which adds information necessary to support transactional record decoding.

The number of **maximum replication slots** must be set to accommodate one open slot per Extract, and in general, no more than one Extract is needed per database. If for example PostgreSQL Native Replication is already in use and is using all of the currently configured replication slots, increase the value to allow for the registration of an Extract.

Maximum write-ahead senders should be set to match the maximum replication slots value.

Optionally, **commit timestamps** can be enabled in the write-ahead log, which when set at the same time logical write-ahead logging is enabled, will track the first DML commit record from that point on, with the correct timestamp value. Otherwise, the first record encountered by Oracle GoldenGate capture will have an incorrect commit timestamp.

After making any of the preceding changes, restart the database.

Topics:

 Database Settings for Azure Database for PostgreSQL, Amazon Aurora, and Amazon RDS

Database Settings for Azure Database for PostgreSQL, Amazon Aurora, and Amazon RDS

Use these instructions to manage the database settings for Azure Database for PostgreSQL, Amazon Aurora, and Amazon RDS.

Azure Database for PostgreSQL

When configuring Oracle GoldenGate for PostgreSQL Capture against an Azure Database for PostgreSQL, **logical decoding** must be enabled and set to LOGICAL.

Read the Microsoft Documentation for instructions.

Other database settings for Azure Database for PostgreSQL can be managed through the **Server parameters** section of the database instance.

For connections to an Azure Database for PostgreSQL instance, the default Azure Connection Security settings require SSL connections. To adhere to this requirement, further steps are required to support SSL connections with Oracle GoldenGate. Follow the content listed under Configuring SSL Support for PostgreSQL for more information.

Amazon Aurora PostgreSQL and Amazon RDS for PostgreSQL

For Amazon Aurora PostgreSQL and Amazon RDS for PostgreSQL, database settings are modified within parameter groups. Review the Amazon AWS documentation for information on how to edit database settings within a new parameter group and assign it to a database instance.

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_WorkingWithParamGroups.html

- Ensure that the database configuration settings listed previously are correct, by verifying them in the parameter group assigned to the instance.
- The wal_level setting for Amazon database services is configured with a parameter called rds.logical_replication, whose default is 0 and should be set to 1 if the database is to be used as source database for Oracle GoldenGate Capture.

Establishing Oracle GoldenGate Credentials

Learn how to create database users for the processes that interact with the database, assign the correct privileges, and secure the credentials from any unauthorized use.

Topics:



- Assigning Credentials to Oracle GoldenGate
- Securing the Oracle GoldenGate Credentials

Assigning Credentials to Oracle GoldenGate

Oracle GoldenGate processes require a database user to capture and deliver data to a PostgreSQL database and it is recommended to create a dedicated PostgreSQL database user for Extract and Replicat.

The following database user privileges are required for Oracle GoldenGate to capture from and apply to a PostgreSQL database.

Extract	Replicat	Purpose
Privileges		
Yes	Yes	Required for database connectivity.
		GRANT CONNECT ON
		DATABASE dbname TO gguser;
Yes	NA	Required for the user to register Extract with a replication slot.
		ALTER USER gguser WITH REPLICATION;
Yes	NA	Required to enable table level supplemental logging (ADD TRANDATA) but can be revoked after TRANDATA is enabled for the table(s).
		ALTER USER gguser WITH SUPERUSER;
		For Azure Database for PostgreSQL, only the Admin user has SUPERUSER authority and is the only user that can enable TRANDATA.
Yes	Yes	For metadata access to tables in the schema to be replicated.
		GRANT USAGE ON SCHEMA tableschema TO gguser;
Yes	Yes	Grant select access on tables to be replicated.
		GRANT SELECT ON ALL TABLES IN SCHEMA tableschema TO
	Yes Yes Yes	Yes Yes Yes NA Yes NA Yes Yes



Privilege	Extract	Replicat	Purpose
INSERT, UPDATE, DELETE, TRUNCATE ON target tables. Alternatively, if replicating every table, then you can use the GRANT INSERT, UPDATE, DELETE, TRUNCATE ON ALL TABLES IN SCHEMA TO to the Replicat user, instead of granting INSERT, UPDATE, DELETE to every table.	NA	Yes	Apply replicated DML to target objects. GRANT INSERT, UPDATE, DELETE, TRUNCATE ON TABLE tablename TO gguser;
Heartbeat and Checkpo	int Table Privileg	es	
CREATE ON DATABASE	Yes	Yes	Required by the Extract and Replicat user to add an Oracle GoldenGate schema for heartbeat and checkpoint table creation. GRANT CREATE ON DATABASE dbname TO gguser; Alternatively, if GGSCHEMA is the same as the user, then the objects can be created under the user by issuing CREATE SCHEMA AUTHORIZATION ggsuser;
CREATE, USAGE ON SCHEMA	Yes	Yes	For heartbeat and checkpoint table creation/deletion if the Extract or Replicat user does not own the objects. GRANT CREATE, USAGE ON SCHEMA ggschema TO gguser;
EXECUTE ON ALL FUNCTIONS	Yes	Yes	For heartbeat update and purge function execution if the user calling the functions does not own the objects. GRANT EXECUTE ON ALL FUNCTIONS IN SCHEMA ggschema TO gguser;



Privilege	Extract	Replicat	Purpose
SELECT, INSERT, Yes UPDATE, DELETE	Yes	For heartbeat and checkpoint table inserts, updates and deletes if the user does not own the objects.	
			GRANT SELECT, INSERT, UPDATE, DELETE, ON ALL TABLES IN SCHEMA ggschema TO gguser;

Securing the Oracle GoldenGate Credentials

To preserve the security of your data, and to monitor Oracle GoldenGate processing accurately, do not permit other users, applications, or processes to log on as, or operate as, an Oracle GoldenGate database user.

Oracle GoldenGate provides different options for securing the log-in credentials assigned to Oracle GoldenGate processes. The recommended option is to use a credential store. You can create one credential store and store it in a shared location where all installations of Oracle GoldenGate can access it, or you can create a separate one on each system where Oracle GoldenGate is installed.

The credential store stores the user name and password for each of the assigned Oracle GoldenGate users. A user ID is associated with one or more aliases, and it is the alias that is supplied in commands and parameter files, not the actual user name or password. The credential file can be partitioned into domains, allowing a standard set of aliases to be used for the processes, while allowing the administrator on each system to manage credentials locally.

See Creating and Populating the Credential Store in *Securing the Oracle GoldenGate Environment* for more information about creating a credential store and adding user credentials.

Configuring a Database Connection

Oracle GoldenGate connects to a PostgreSQL database through an ODBC (Open Database Connectivity) driver and requires a system Data Source Name (DSN) be created with the correct database connection details for each source and target PostgreSQL database.

This section contains instructions for setting up the DSN connections that Extract and Replicat will use.

Ensure that you have installed and configured the driver prior to creating a DSN, by following the Installing the DataDirect driver for PostgreSQL instructions in *Installing Oracle GoldenGate*.

Topics:

- Configuring a Database Connection in Linux
- Configuring SSL Support for PostgreSQL



Configuring a Database Connection in Linux

To create a database connection in Linux, set up a data source name (DSN) inside the /etc/odbc.ini file.

1. Create a DSN for each source or target database in the /etc/odbc.ini file.

```
sudo vi /etc/odbc.ini
#Sample DSN entries
[ODBC Data Sources]
PG src=DataDirect 7.1 PostgreSQL Wire Protocol
PG tgt=DataDirect 7.1 PostgreSQL Wire Protocol
[ODBC]
IANAAppCodePage=4
InstallDir=/u01/app/ogg
[PG src]
Driver=/u01/app/ogg/lib/GGpsql25.so
Description=DataDirect 7.1 PostgreSQL Wire Protocol
Database=sourcedb
HostName=remotehost
PortNumber=5432
PG tgt]
Driver=/u01/app/ogg/lib/GGpsgl25.so
Description=DataDirect 7.1 PostgreSQL Wire Protocol
Database=targetdb
HostName=remotehost
PortNumber=5432
```

In the preceding examples:

PG_src and PG_tgt are user defined names of a source and target database DSN that will be referenced by Oracle GoldenGate processes, such as Extract or Replicat. DSN names are allowed up to 32 alpha-numeric characters in length, excluding special keyboard characters except for the underscore and dash.

IANAAppCodePage=4 is the default setting but can be modified according to the following guidance, when the database character set is not Unicode.

https://docs.progress.com/bundle/datadirect-connect-odbc-71/page/IANAAppCodePage_9.html#IANAAppCodePage_9

InstallDir is the location of the Oracle GoldenGate installation folder.

Driver is the location of the Oracle GoldenGate installation home, <code>\$OGG_HOME/lib/GGpsg125.so</code> file.

Database is the name of the source or target database.

HostName is the database host IP address or host name.

PortNumber is the listening port of the database.

You can also provide a LogonID and Password for the Extract or Replicat user, but these will be stored in clear text and it is recommended instead to leave these fields out of the DSN and instead store them in the Oracle GoldenGate wallet as a credential alias, and reference them with the USERIDALIAS parameter in Extract and Replicat.

2. Save and close the odbc.ini file.

Configuring SSL Support for PostgreSQL

SSL can be enabled by configuring the PostgreSQL configuration file (\$PGDATA/postgresql.conf). For details, see Configuring SSL Support (PostgreSQL) in the Securing the Oracle GoldenGate Environment.



Azure Database for PostgreSQL defaults to enforce SSL connections. To adhere to this requirement, perform the requirements listed here, or optionally, you can disable enforcing SSL connections from the **Connection security** settings of the database instance using the Microsoft Azure Portal.

Preparing Tables for Processing

The following table attributes must be addressed in an Oracle GoldenGate environment for PostgreSQL.

Topics:

- Disabling Triggers and Cascade Constraints on the Target
- Ensuring Row Uniqueness for Tables
- Enabling Table-Level Supplemental Logging

Disabling Triggers and Cascade Constraints on the Target

If Oracle GoldenGate is configured to capture DML operations from source tables that occur due to trigger operations or cascade constraints, then disable the triggers and cascade delete and cascade update constraints on the target tables.

If not disabled, the same trigger or constraint gets activated on the target table and becomes redundant because of the replicated data. Consider the following example, where the source tables are <code>emp_src</code> and <code>salary_src</code> and the target tables are <code>emp_targ</code> and <code>salary_targ</code>

- 1. A delete is issued for emp src.
- 2. It cascades a delete to salary src.
- 3. Oracle GoldenGate sends both deletes to the target.
- 4. The parent delete arrives first and is applied to emp targ.
- The parent delete cascades a delete to salary targ.
- 6. The cascaded delete from salary src is applied to salary targ.
- The row cannot be located because it was already deleted in step 5.



In the Replicat MAP statements, map the source tables to appropriate targets, and map the child tables that the source tables reference with triggers or foreign-key cascade constraints. Triggered and cascaded child operations must be mapped to appropriate targets to preserve data integrity. Include the same parent and child source tables in the Extract TABLE parameters.

Ensuring Row Uniqueness for Tables

Oracle GoldenGate requires some form of unique row identifier on the source and target tables to locate the correct target rows for replicated updates and deletes.

Unless a KEYCOLS clause is used in the TABLE or MAP statement, Oracle GoldenGate selects a row identifier to use in the following order of priority:

- Primary key
- 2. First unique key alphanumerically that does not contain a timestamp or non-materialized computed column.
- 3. If none of the preceding key types exist (even though there might be other types of keys defined on the table) Oracle GoldenGate constructs a pseudo key of all columns that the database allows to be used in a unique key, excluding those that are not supported by Oracle GoldenGate in a key or those that are excluded from the Oracle GoldenGate configuration. For PostgreSQL LOB types such as text, xml, bytea, char, varchar, Oracle GoldenGate supports these columns as a primary key in source or target tables up to a length of 8191 bytes.

Note:

If there are other, non-usable keys on a table or if there are no keys at all on the table, Oracle GoldenGate logs an appropriate message to the report file. Constructing a key from all of the columns impedes the performance of Oracle GoldenGate on the source system. On the target, this key causes Replicat to use a larger, less efficient WHERE clause.

For tables that have no uniqueness and have repeat rows with the same values, Replicat will Abend on update and delete operations for these rows.

4. If a table does not have an appropriate key, or if you prefer that the existing key(s) are not used, you can define a substitute key, if the table has columns that always contain unique values. You define this substitute key by including a KEYCOLS clause within the Extract TABLE parameter and the Replicat MAP parameter. The specified key will override any existing primary or unique key that Oracle GoldenGate finds. See TABLE | MAP in Reference for Oracle GoldenGate.

Enabling Table-Level Supplemental Logging

Enabling Supplemental logging is a process in which Oracle GoldenGate sets source database table level logging to support change data capture of source DML operations, and depending on the level of logging, to include additional, unchanged columns which would be needed in cases such as bi-directional replication with conflict detection and resolution configured.

There are four levels of table level logging in PostgreSQL, which equate to the REPLICA IDENTITY setting of a table, and those include NOTHING, USING INDEX, DEFAULT, and FULL.

Oracle GoldenGate requires FULL logging for use cases that require uncompressed trail records and Conflict Detection and Resolution, but in cases where tables have a Primary Key

or Unique Index whose changes are being replicated in a simple uni-directional configuration or where full before-images or uncompressed records are not needed, then the DEFAULT level is acceptable. NOTHING and USING INDEX logging levels are not supported by Oracle GoldenGate and cannot be set with ADD TRANDATA.

The following is the syntax for issuing ADD TRANDATA from GGSCI.

GGSCI> DBLOGIN SOURCEDB dsn_name USERIDALIAS alias_name GGSCI> ADD TRANDATA schema.tablename ALLCOLS



For tables that have a primary key or unique index, the ALLCOLS option is required in order to set FULL logging for the table, otherwise DEFAULT logging is set.

 ${ t FULL}$ logging is always set for tables without a primary key or unique index, regardless of whether ${ t ALLCOLS}$ is specified or not.

To check the level of supplemental logging:

GGSCI> INFO TRANDATA schema.tablename



17

Configuring Extract

This chapter contains instructions for configuring the Oracle GoldenGate capture process to capture initial load and transactional data from a PostgreSQL database.

Topics:

- About Extract
- · Prerequisites for Creating an Extract
- Creating a Change Data Capture Extract
- · Modifying DDL for Extract

About Extract

For Oracle GoldenGate for PostgreSQL, there are two types of Extracts that can be created.

- Initial Load Extract
- Change Data Capture Extract
- Extract Deployment Options

Initial Load Extract

An Initial Load Extract is used to read all records from a table and write them to an EXTFILE or RMTFILE. Initial load Extracts are created with the SOURCEISTABLE option of the ADD EXTRACT command and do not maintain checkpointing for recovery.

For more information on the Initial Load process, see Instantiating Oracle GoldenGate with an Initial Load in *Administering Oracle GoldenGate*.

Change Data Capture Extract

A Change Data Capture Extract is used to capture transactional data changes from that point in time at which it is created or positioned into the write-ahead log.

The Oracle GoldenGate Extract process for PostgreSQL receives logical records from the PostgreSQL test_decoding database plugin and writes them in commit order into trail files for downstream consumption by a Replicat.

Extract Deployment Options

- Local deployment: For a local deployment, the source database and Oracle GoldenGate are installed on the same server. No extra consideration is needed for local deployments.
- Remote deployment: For a remote deployment, the source database and Oracle GoldenGate are installed on separate servers. Remote deployments are the only option available for supporting cloud databases, such as Azure for PostgreSQL or Amazon Aurora PostgreSQL.

For remote deployments, operating system endianness between the database server and Oracle GoldenGate server need to be the same.

Server time and time zones of the Oracle GoldenGate server should be synchronized with that of the database server. If this is not possible, then positioning of an Extract when creating or altering one will need to be done by LSN.

In remote capture use cases, using SQLEXEC may introduce additional latency, as the SQLEXEC operation must be done serially for each record that the Extract processes. If special filtering that would require a SQLEXEC is done by a remote hub Extract and the performance impact is too severe, it may become necessary to move the Extract process closer to the source database.

With remote deployments, low network latency is important, and it is recommended that the network latency between the Oracle GoldenGate server and the source database server be less than 1 millisecond.

Prerequisites for Creating an Extract

Review the Installing Oracle GoldenGate for PostgreSQL and ensure that the DataDirect drivers are installed correctly, which varies depending on the operating system.

Ensure that the PostgreSQL Client Authentication Configuration file, \$PGDATA/pg_hba.conf, on the database server is configured to allow connections from the Oracle GoldenGate server, if installed remotely. See https://www.postgresql.org/docs/13/auth-pg-hba-conf.html for more information.

Registering an Extract for PostgreSQL

Registering an Extract for PostgreSQL

An Extract for PostgreSQL must be registered with the database and be granted a reserved replication slot. Replication slots are allocated through the database configuration setting max_replication_slots and can be configured as discussed in the *Database Configuration* topic of this document.

Follow these instructions to register an Extract. Extract registration must be done prior to creating an Extract. See REGISTER EXTRACT in the *Command Line Interface Reference for Oracle GoldenGate* guide for more information.

1. Using GGSCI, connect to the DSN for the source database.

```
GGSCI> DBLOGIN SOURCEDB dsn USERIDALIAS alias
```

2. Register the Extract using the GGSCI command. This command internally creates the replication slot. Extract names cannot be more than 8 alpha-numeric characters.

```
GGSCI> REGISTER EXTRACT extname
```

Creating a Change Data Capture Extract

These steps configure a CDC Extract to capture transactional data from a source PostgreSQL database.

Note:

One Extract per database is generally sufficient, but multiple Extracts are allowed if the replication slots are available.

1. In GGSCI, create the Extract parameter file.

EDIT PARAMS extname

In this sample, <code>extname</code> is the name of the primary Extract and matches the name of the Extract that was registered with the database in the previous steps.

2. Enter the Extract parameters in the order shown, starting a new line for each parameter statement.

Sample basic parameters for Extract:

EXTRACT extname

SOURCEDB dsn_name USERIDALIAS alias

EXTTRAIL ./dirdat/ep

GETTRUNCATES

TABLE schema.object;

Parameter	Description
EXTRACT extname	extname is the name of the Extract and cannot be more than 8 alpha-numeric characters in length. For more information, see extract in Reference for Oracle GoldenGate.
SOURCEDB dsn_name	Specifies the name of the database connection DSN.
USERIDALIAS alias	Specifies the alias of the database login credential of the user that is assigned to Extract. This credential must exist in the Oracle GoldenGate credential store. For more information, see Establishing Oracle GoldenGate Credentials.
EXTTRAIL trailname	Specifies a two character, local trail to which the primary Extract writes captured data.
GETTRUNCATES	Optional parameter but needed in order to capture truncation operations.



Description
Specifies the database object for which to
capture data.
 TABLE specifies a table or a wildcarded set of tables.
 schema is the schema name or a wildcarded set of schemas.
 object is the table or sequence name, or a wildcarded set of those objects.
 * is a wildcard for all tables in the schema.
Terminate the parameter statement with a semi- colon.
To exclude a name from a wildcard specification, use the SCHEMAEXCLUDE, TABLEEXCLUDE, and EXCLUDEWILDCARDOBJECTSONLY parameters as appropriate.



If the schema of tables to be captured from is the same as the schema in GGSCHEMA of the GLOBALS file, which is not recommended, then you cannot use schema.* in the TABLE statement.

- Enter any optional Extract parameters that are recommended for your configuration. You
 can edit this file at any point before starting processing by using the EDIT PARAMS
 command.
- 4. Save and close the file.
- 5. Add the Extract and its associated trail file.

```
GGSCI> ADD EXTRACT extname, TRANLOG, BEGIN NOW GGSCI> ADD EXTTRAIL ./dirdat/ep, EXTRACT extname
```

6. Start the Extract.

Modifying DDL for Extract

Follow these steps to modify DDL for Extract for PostgreSQL:

- Pause or stop the application data to the table(s), which need to be modified.
- 2. Ensure that Extract processes all the transactions prior to making any DDL changes. An Event Marker table helps ensure full completion of transactions.
- **3.** Stop the Extract.
- **4.** At source, execute DELTE TRANDATA for the table(s) on which DDL modification has to be performed.
- 5. Run the ALTER TABLE statement to add or drop the column in or from the table(s).
- 6. Enable trandata (ADD TRANDATA) for the same table(s) at source.



7. Start Extract.



18

Configuring Replicat

This chapter contains instructions for configuring the Replicat apply process to deliver data to a target PostgreSQL database.

Topics:

- About Replicat
- · Prerequisites for Creating a Replicat
- · Creating a Replicat

About Replicat

The Oracle GoldenGate Replicat for PostgreSQL reads data from Oracle GoldenGate source trail files and delivers the data to a target PostgreSQL database. The source trail data can be from any database that Oracle GoldenGate capture supports.

Available Replicats for PostgreSQL are Classic and Coordinated.

For more information on the differences between types of Replicats, review the Creating an Online Replicat Group content in the *Administering Oracle GoldenGate* guide.

Replicat Deployment Options

Replicat Deployment Options

- Local deployment: For a local deployment, the target database and Oracle GoldenGate are installed on the same server. No extra consideration is needed for local deployments.
- Remote deployment: For a remote deployment, the target database and Oracle GoldenGate are installed on separate servers. Remote deployments are the only option available for supporting cloud databases, such as Azure for PostgreSQL or Amazon Aurora PostgreSQL.

For remote deployments, operating system endianness between the database server and Oracle GoldenGate server needs to be the same.

With remote deployments, low network latency is important, and it is recommended that the network latency between the Oracle GoldenGate server and the target database server be less than 1 millisecond.

Prerequisites for Creating a Replicat

Review the Installing the DataDirect driver for PostgreSQL in *Installing Oracle GoldenGate* and ensure that the DataDirect drivers are installed correctly, which varies depending on the operating system.

Ensure that the PostgreSQL Client Authentication Configuration file, <code>\$PGDATA/pg_hba.conf</code>, on the database server is configured to allow connections from the Oracle GoldenGate server, if installed remotely. See https://www.postgresql.org/docs/13/auth-pg-hba-conf.html for more information.

Creating a Checkpoint Table

Creating a Checkpoint Table

A checkpoint table is used by a Replicat in the target database for recovery positioning when restarting a Replicat. A checkpoint table is optional (but recommended) for a Classic Replicat and required for a Coordinated Replicat.

The checkpoint table needs to be created under an existing schema in the database and by default will attempt to create the table in the schema listed in the <code>GLOBALS</code> file, <code>GGSCHEMA</code> parameter. Ensure that the schema listed in <code>GLOBALS</code> exists in the database and that the Replicat user has permissions to use the schema and create tables in it.

These steps demonstrate creating a checkpoint table for a Classic and Coordinated Replicat.

Using GGSCI, connect to the DSN for the target database.

```
GGSCI> DBLOGIN SOURCEDB dsn USERIDALIAS alias
```

2. Add the checkpoint table using the GGSCI command.

```
GGSCI> ADD CHECKPOINTTABLE ggadmin.oggcheck
```

Creating a Replicat

These steps create a Replicat to deliver transactional data to a target PostgreSQL database.

1. In GGSCI, create the Replicat parameter file.

```
EDIT PARAMS repnm
```

In this sample, repnm is a name of the Replicat. For Classic Replicat, the name can be no more than 8 alpha-numeric characters in length. For Coordinated Replicat, the name must be five or less alpha-numeric characters in length.

Enter the Replicat parameters in the order shown, starting a new line for each parameter statement.

Sample basic parameters for Classic Replicat:

```
REPLICAT repnm

TARGETDB dsn_name USERIDALIAS alias

BATCHSQL

GETTRUNCATES

MAP schema.object, TARGET schema.object;
```

Sample basic parameters for Coordinated Replicat:

```
REPLICAT repnm

TARGETDB dsn_name USERIDALIAS alias

BATCHSQL

GETTRUNCATES

MAP schema.object1, TARGET schema.object1, THREAD (1);

MAP schema.object2, TARGET schema.object2, THREAD (2);

MAP schema.object3, TARGET schema.object3, THREAD (3);
```



Parameter	Description	
REPLICAT repnm	repnm is the name of the Replicat and cannot be more than 8 alpha-numeric characters in length for Classic Replicat and 5 or less for Coordinated Replicat. For more information, see REPLICAT in Reference for Oracle GoldenGate.	
TARGETDB dsn_name	Specifies the name of the database connection DSN.	
USERIDALIAS alias	Specifies the alias of the database login credential of the user that is assigned to Replicat. This credential must exist in the Oracle GoldenGate credential store. For more information, see Establishing Oracle GoldenGate Credentials.	
BATCHSQL GETTRUNCATES	Optional parameters for Replicat that supports transaction batching and replication of truncate operations.	
MAP schema.object, TARGET schema.object; or MAP schema.*, TARGET schema.*;	 Specifies the relationship between a source table and the corresponding target object or objects. MAP specifies the source table or a wildcarded set of tables. TARGET specifies the target table or a wildcarded set of tables. schema is the schema name or a wildcarded set of schemas. object is the name of a table or a wildcarded set of tables. THREAD assigns table operations to a specific coordinated Replicat thread. Terminate the parameter statement with a semi-colon. To exclude objects from a wildcard specification, use the MAPEXCLUDE parameter. For more information and for additional options that control data filtering, mapping, and manipulation, see MAP in Reference for Oracle GoldenGate. 	

- 3. Enter any optional Replicat parameters that are recommended for your configuration. You can edit this file at any point before starting processing by using the EDIT PARAMS command.
- 4. Save and close the file.
- 5. Add the Replicat, which in this example, will be a Coordinated Replicat.

GGSCI> ADD REPLICAT repnm, COORDINATED, EXTTRAIL ./dirdat/ep, CHECKPOINTTABLE ggadmin.oggcheck

6. Start the Replicat.



Additional Considerations

- Adding a Heartbeat Table
- Enabling Bi-Directional Loop Detection
- Deleting an Extract
- Removing Table-level Supplemental Logging

Adding a Heartbeat Table

Oracle GoldenGate for PostgreSQL supports a heartbeat table configuration, with some limitations regarding the update and purge tasks, which will be pointed out later. Adding a heartbeat table to both the source and target systems is optional but is useful in determining latency issues and to which process in the replication stream such issues may be occurring.

To add a heartbeat table for a database, review the required privileges in the Database Privileges for PostgreSQL and ensure that the correct database privileges are assigned to the Extract or Replicat user.

A schema in the database needs to be created and this should match the name of the schema used for the GGSCHEMA parameter of the GLOBALS file. The schema can be a unique schema that is not the same as the Extract or Replicat user, or can be the same as that user, but should always be reserved for Oracle GoldenGate objects only and should not be part of the user table schemas being replicated.

Using GGSCI, connect to the DSN for the source and target databases.

```
GGSCI> DBLOGIN SOURCEDB dsn USERIDALIAS alias
```

2. Add the heartbeat table using the GGSCI command.

```
GGSCI> ADD HEARTBEATTABLE
```

Optionally, for a target only database, one that is used for unidirectional replication only, you can include the TARGETONLY option, which will not create a heartbeat record update function.

```
GGSCI> ADD HEARTBEATTABLE TARGETONLY
```

To learn about heartbeat update and purge functions, see:

Running the Heartbeat Update and Purge Function

Running the Heartbeat Update and Purge Function

The heartbeat table and associated functions are created from the ADD HEARTBEATTABLE command, however for PostgreSQL, there is no automatic scheduler to call the functions.

One main function controls both the heartbeat record update and the heartbeat history table purge functions. The default settings for both of these features are 60 seconds for the update frequency and 1 day for the history record purge, which deletes all records older than 30 days by default.

To call the main heartbeat record function, users should create an operating system level job that executes "select ggschema.gg_hb_job_run();". When this function is called, it will take into account the update frequency settings and history record purge settings and use those values regardless of the scheduler settings for the function call.

For example, users can create a **Cron Job** with the following syntax, and have it run every minute.

```
*****PGPASSWORD="gguserpasswd" psql -U gguser -d dbname -h remotehost -p 5432 -c "select ggschema.gg_hb_job_run();" >/dev/null 2>&1
```

pgAdmin, or pg_cron are other programs that could be used to schedule the function call.

Enabling Bi-Directional Loop Detection

Loop detection is a requirement for bi-directional and multi-directional implementations of Oracle GoldenGate, so that an Extract for one source database does not recapture transactions sent by a Replicat from another source database.

With the CDC Extract capture method, by default, any transaction committed by a Replicat into a database where an Extract is configured, will recapture that transaction from the Replicat as long as supplemental logging is enabled for those tables that the Replicat is delivering to.

In order to ignore recapturing transactions that are applied by a Replicat, you must use the TRANLOGOPTIONS FILTERTABLE parameter for the CDC Extract. The table used as the filtering table will be the Oracle GoldenGate checkpoint table that you must create for the Replicat.

To create a Filter Table and enable Supplemental Logging:

 On each source database, ensure that a checkpoint table for use by Replicats has been created. For example:

```
ADD CHECKPOINTTABLE ggadmin.oggcheck
```

It is recommended that you use the same schema name as used in the GGSCHEMA parameter of the GLOBALS file.

2. Enable supplemental logging for the checkpoint table. For example:

```
ADD TRANDATA ggadmin.oggcheck ALLCOLS
```

3. Ensure that the Replicat is created with the checkpoint table information.

```
ADD REPLICAT reptgt1, EXTTRAIL ./dirdat/e2, CHECKPOINTTABLE ggadmin.oggcheck
```



4. Configure each Extract with the IGNOREREPLICATES (on by default) and FILTERTABLE parameters, using the Replicat's checkpoint table for the filtering table.

```
TRANLOGOPTIONS IGNOREREPLICATES
TRANLOGOPTIONS FILTERTABLE ggadmin.oggcheck
```



Oracle GoldenGate for PostgreSQL supports only one FILTERTABLE statement per Extract, so for multi-directional implementations, ensure each Replicat uses the same checkpoint table in the database that they deliver to.

Deleting an Extract

When removing an individual Extract from use against a source PostgreSQL database, or uninstalling Oracle GoldenGate, the Extract that was registered with a replication slot in the database must be unregistered in order to remove the replication slot entry, otherwise an ever-increasing database log size can occur.

Perform the following steps to remove and unregister the Extract when no longer needed.

1. Using GGSCI, connect to the DSN for the source and target databases.

```
GGSCI> DBLOGIN SOURCEDB dsn USERIDALIAS alias
```

Delete the Extract first.

```
GGSCI> DELETE EXTRACT extname
```

3. Unregister the Extract.

GGSCI> UNREGISTER EXTRACT extname

Removing Table-level Supplemental Logging

If a table is no longer required to be captured by Oracle GoldenGate and the TABLE parameter for the table has been removed from the Extract parameter file, or TABLEEXCLUDE is used to exclude the table from a wildcard statement, then supplemental logging can be removed from the table.



If the Extract resolves a table that does not have supplemental logging enabled, it will Abend depending on the type of DML operation.

Using DELETE TRANDATA to remove supplemental logging sets the Replicat Identity level of the table to NOTHING. Supplemental logging can be disabled using the DELETE TRANDATA command within GGSCI.

The following is the syntax for issuing DELETE TRANDATA from GGSCI.

GGSCI> DBLOGIN SOURCEDB dsn_name USERIDALIAS alias_name GGSCI> DELETE TRANDATA schema.tablename

To check the level of supplemental logging:

GGSCI> INFO TRANDATA schema.tablename



Part VII

Using Oracle GoldenGate for SQL Server

With Oracle GoldenGate for SQL Server, you can perform initial loads and capture transactional data from supported SQL Server versions and replicate the data to a SQL Server database or other supported Oracle GoldenGate targets, such as an Oracle Database.

Oracle GoldenGate for SQL Server supports data filtering, mapping, and transformations unless noted otherwise in this documentation.

This part describes tasks for configuring and running Oracle GoldenGate for SQL Server.

Topics:

- Understanding What's Supported for SQL Server
 This chapter contains information on database and table features supported by Oracle GoldenGate for SQL Server.
- Preparing the System for Oracle GoldenGate
- Preparing the Database for Oracle GoldenGate CDC Capture
 Learn how to enable supplemental logging in the source database tables that are to be used for capture by the Extract for SQL Server and how to purge older CDC staging data.
- Requirements Summary for Capture and Delivery of Databases in an Always On Availability Group
- CDC Capture Method Operational Considerations
 This section provides information about the SQL Server CDC Capture options, features, and recommended settings.



20

Understanding What's Supported for SQL Server

This chapter contains information on database and table features supported by Oracle GoldenGate for SQL Server.

Topics:

- Supported Objects and Operations for SQL Server
- Non-Supported Objects and Operations for SQL Server
- Supported SQL Server Data Types
- System Schemas for SQL Server
- Non-Supported SQL Server Data Types and Features

Supported Objects and Operations for SQL Server

The following objects and operations are supported:

- Oracle GoldenGate supports capture of transactional DML from user tables and delivery to user tables and writeable views.
- TEXT, NTEXT, IMAGE, VARBINARY, VARBINARY (MAX) VARCHAR (MAX), and NVARCHAR (MAX) columns are supported in their full size.
- Oracle GoldenGate supports the maximum row sizes that are permitted for tables that are enabled for SQL Server Change Data Capture.
- Oracle GoldenGate supports capture from tables enabled with PAGE and ROW compression.
 For partitioned tables that use compression, all partitions must be enabled with the same compression type.
- Oracle GoldenGate supports capture for partitioned tables if the table has the same physical layout across all partitions.
- The sum of all column lengths for a table to be captured from must not exceed the length that SQL Server allows for enabling Change Data Capture for the table. If the sum of all column lengths exceeds what is allowed by the SQL Server procedure sys.sp.cdc_enable_table, then ADD TRANDATA cannot be added for that table. The maximum allowable record length decreases as more columns are present, so there is an inverse relationship between maximum record length and the number of columns in the table.

Non-Supported Objects and Operations for SQL Server

The following objects and operations are not supported:

- Parallel Replicat is supported with Oracle GoldenGate for SQL Server.
- For source databases, operations that are not supported by SQL Server Change Data Capture, such as TRUNCATE statements. Refer to Microsoft SQL Server Documentation for

- a complete list of the operations that are limited by enabling SQL Server Change Data Capture.
- Oracle GoldenGate for SQL Server does not support the capture or delivery of DDL changes for SQL Server and extra steps are required for Oracle GoldenGate processes on the source and target to handle any table level DDL changes, including table index rebuild operations. See Requirements for Table Level DDL Changes.
- Views are not supported.
- Operations by the TextCopy utility and WRITETEXT and UPDATETEXT statements. These
 features perform operations that either are not logged by the database or are only partially
 logged, so they cannot be supported by the Extract process.
- Partitioned tables that have more than one physical layout across partitions.
- Partition switches against a source table. SQL Server Change Data Capture treats partition switches as DDL operations, and the data moved from one partition to another is not logged in the CDC tables, so you must follow the procedures in Requirements for Table Level DDL Changes to manually implement a partition switch when the table is enabled for supplemental logging.
- Due to a limitation with SQL Server's Change Data Capture, column level collations that
 are different from the database collation, may cause incorrect data to be written to the
 CDC tables for character data and Extract will capture them as they are written to the CDC
 tables. It is recommended that you use NVARCHAR, NCHAR or NTEXT data type for columns
 containing non-ASCII data or use the same collation for table columns as the database.
 For more information see, About Change Data Capture (SQL Server).
- Due to a limitation with SQL Server's Change Data Capture, NOOPUPDATES are not captured by the SQL Server Change Data Capture agent so there are no records for Extract to capture for no-op update operations.
- Requirements for Table Level DDL Changes

Requirements for Table Level DDL Changes

Oracle GoldenGate for SQL Server, including table index rebuild operations, does not support the capture or delivery of DDL changes for SQL Server. Extra steps are required for Oracle GoldenGate processes on the source and target to handle any table-level DDL changes.

Operations considered to be table level DDL changes include, but are not limited to ALTER TABLE, TRUNCATE, index rebuilds, and partition switches.

To avoid data inconsistencies due to table level DDL changes, the following steps are required.

- 1. Source: Pause or Stop application data to the table or tables to be modified.
- 2. Source: Ensure that there are no open transactions against the table to be modified.
- 3. Source: Ensure that the SQL Server CDC Capture job processes all remaining transactions for the table that is to be modified.
- 4. Source: Ensure that the Extract processes all the transactions for the table that is to be modified, prior to making any DDL changes.
- 5. Target: Ensure that the Replicat processes all the transactions for the table that is to be modified, prior to making any DDL changes.
- 6. Optionally, implementing an Event Marker table can be used to determine when all of the remaining transactions have been processed for the table that is to be modified, and handle the coordination of when to correctly stop the Extract and Replicat.



- Source: Stop the Extract process.
- Target: Stop the Replicat process.
- Source: Disable supplemental logging for the table to be modified by running DELETE TRANDATA.
- 10. Source: Make table DDL changes to the source table.
- 11. Target: Make table DDL changes to the target table.
- **12.** Source: Re-enable supplemental logging by running ADD TRANDATA to the table(s) after the modifications have been performed.
- 13. Source: Start the Extract.
- 14. Target: Start the Replicat.
- 15. Source: Resume application data to the table or tables that were modified.

Supported SQL Server Data Types

The following data types are supported for capture and delivery, unless specifically noted in the limitations that follow:

- Binary Data Types
 - (binary, varbinary, varbinary (max))
 - (varbinary (max) with FILESTREAM)
- Character Data Types
 - (char, nchar, nvarchar, nvarchar (max), varchar, varchar (max))
- Date and Time Data Types
 - (date, datetime2, datetime, datetimeoffset, smalldatetime, time)
- Numeric Data Types
 - (bigint, bit, decimal, float, int, money, numeric, real, smallint, smallmoney, tinyint)
- LOBs
 - (image, ntext, text)
- Other Data Types
 - (timestamp, uniqueidentifier, hierarchyid, geography, geometry, sql variant (Delivery only), XML)
- Oracle GoldenGate for SQL Server can replicate column data that contains SPARSE settings..

Limitations:

- Oracle GoldenGate does not support filtering, column mapping, or manipulating large objects larger than 4KB. Full Oracle GoldenGate functionality can be used for objects of up to 4KB.
- Oracle GoldenGate treats XML data as a large object (LOB), as does SQL Server when the XML does not fit into a row. SQL Server extended XML enhancements (such as lax validation, DATETIME, union functionality) are not supported.

- A system-assigned TIMESTAMP column or a non-materialized computed column cannot be part of a key. A table containing a TIMESTAMP column must have a key, which can be a primary key or unique constraint, or a substitute key specified with a KEYCOLS clause in the TABLE or MAP statements. For more information see Assigning Row Identifiers.
- Oracle GoldenGate supports multibyte character data types and multi byte data stored in character columns. Multibyte data is supported only in a like-to-like, SQL Server configuration. Transformation, filtering, and other types of manipulation are not supported for multibyte character data.
- If capture of data for TEXT, NTEXT, IMAGE, VARCHAR (MAX), NVARCHAR (MAX) and VARBINARY (MAX) columns will exceed the SQL Server default size set for the max text repl size option, extend the size. Use sp_configure to view the current value of max text repl size and adjust the option as needed.
- Oracle GoldenGate supports UDT and UDA data of up to 2 GB in size. All UDTs except SQL_Variant are supported.
- Common Language Runtime (CLR), including SQL Server built-in CLR data types (such as, geometry, geography, and hierarchy ID), are supported. CLR data types are supported only in a like-to-like SQL Server configuration. Transformation, filtering, and other types of manipulation are not supported for CLR data.
- VARBINARY (MAX) columns with the FILESTREAM attribute are supported up to a size of 4 GB. Extract uses standard Win32 file functions to read the FILESTREAM file.
- The range and precision of floating-point numbers depends on the host machine. In general, precision is accurate to 16 significant digits, but you should review the database documentation to determine the expected approximations. Oracle GoldenGate rounds or truncates values that exceed the supported precision.
- Oracle GoldenGate supports time stamp data from 0001/01/03:00:00:00:00 to 9999/12/31:23:59:59. If a time stamp is converted from GMT to local time, these limits also apply to the resulting time stamp. Depending on the time zone, conversion may add or subtract hours, which can cause the time stamp to exceed the lower or upper supported limit.

Limitations on Computed Columns:

- Computed columns, either persisted or non-persisted, are not supported by Microsoft's Change Data Capture. Therefore, no data is written to the trail for columns that contain computed columns. To replicate data for non-persisted computed columns, use the FETCHCOLS or FETCHMODCOLS option of the TABLE parameter to fetch the column data from the table. Keep in mind that there can be discrepancies caused by differences in data values between the time that the column was changed in the data base and the time that Extract fetches the data for the transaction record that is being processed.
- Replicat does not apply DML to any computed column, even if the data for that column is
 in the trail, because the database does not permit DML on that type of column. Data from a
 source persisted computed column, or from a fetched non- persisted column, can be
 applied to a target column that is not a computed column.
- In an initial load, all of the data is selected directly from the source tables, not the transaction log. Therefore, in an initial load, data values for all columns, including non-persisted computed columns, is written to the trail or sent to the target, depending on the method that is used. As when applying change data, however, Replicat does not apply initial load data to computed columns, because the database does not permit DML on that type of column.
- Oracle GoldenGate does not permit a non-persisted computed column to be used in a KEYCOLS clause in a TABLE or MAP statement.



- If a unique key includes a non-persisted computed column and Oracle GoldenGate must
 use the key, the non-persisted computed column is ignored. This may affect data integrity if
 the remaining columns do not enforce uniqueness.
- If a unique index is defined on any non-persisted computed columns, it is not used.
- If a unique key or index contains a non-persisted computed column and is the only unique identifier in a table, Oracle GoldenGate must use all of the columns as an identifier to find target rows. Because a non-persisted computed column cannot be used in this identifier, Replicat may apply operations containing this identifier to the wrong target rows.

System Schemas for SQL Server

The following schemas or objects are not be automatically replicated by Oracle GoldenGate unless they are explicitly specified without a wildcard.

- "sys"
- "cdc"
- "INFORMATION SCHEMA"
- "quest"

Non-Supported SQL Server Data Types and Features

- SQL Variant data type is not supported for capture.
- Tables that contain unsupported data types may cause Extract to Abend. As a workaround, you must remove TRANDATA from those tables and remove them from the Extract's TABLE statement, or use the Extract's TABLEEXCLUDE parameter for the table.



Preparing the System for Oracle GoldenGate

This chapter contains steps to take so that the database with which Oracle GoldenGate interacts is correctly configured to support Oracle GoldenGate capture and delivery. Some steps apply only to a source system, some only to a target, and some to both. **Topics:**

- Prepare Database Users and Privileges
- Configuring a Database Connection
- Preparing Tables for Processing
- Globalization Support

Prepare Database Users and Privileges

Learn about required database users, privileges, and permissions for Oracle GoldenGate for SQL Server, including supported SQL Server cloud databases.

- Oracle GoldenGate for SQL Server
- Amazon RDS User Permissions and Requirements
- Azure SQL Database

Oracle GoldenGate for SQL Server

Oracle GoldenGate processes require a database user in order to capture from and apply data to a SQL Server database and it is recommended to create a dedicated database user to be used exclusively by Oracle GoldenGate processes.

Oracle GoldenGate for SQL Server supports SQL Server authentication for all of its certified platforms and Windows authentication for Classic Architecture only when Oracle GoldenGate is installed on a Windows server.

- To use Windows authentication for Oracle GoldenGate Classic Architecture, the Extract
 and Replicat processes inherit the login credentials of the Manager process. By default,
 the Manager process runs interactively as the user logged on to the Windows server or
 optionally can be added as a Windows Service with a default service name of GGSMGR.
 Whichever method that the Manager process is using, the user that it is running as needs
 to have the required SQL Server privileges listed above.
- To use SQL Server authentication, create a dedicated SQL Server login for Extract and Replicat and assign the privileges listed below.

SQL Server and Azure SQL Managed Instance

The following user requirements and minimum database privileges and permissions are required for Oracle GoldenGate to capture from and apply to a SQL Server or Azure SQL Managed Instance database.

 Create a dedicated login for Oracle GoldenGate for SQL Server or Azure SQL Managed Instance.

- Add the login as a user to the msdb database and to the source or target database.
- Create a schema in the source or target database, to be used for objects required for Oracle GoldenGate. This schema should map to the GGSCHEMA value used in the GLOBALS parameter file.
- **4.** Enable the following privileges and permissions for the Oracle GoldenGate user based on whether the user is for an Extract, or for a Replicat.

Table 21-1 Privileges and Permissions for Oracle GoldenGate User

Privilege	Extract	Replicat	Syntax			
msdb Database Roles and Privileges						
SQLAgentReaderRole	Yes	No	ALTER ROLE SQLAgentReaderRole ADD MEMBER gguser;			
SQLAgentUserRole	Inherited	Yes	ALTER ROLE SQLAgentUserRole ADD MEMBER gguser;			
SELECT ON sysjobactivity	Yes	No	GRANT SELECT ON msdb.dbo.sysjobactivity TO gguser;			
SELECT ON sysjobs	Yes	No	GRANT SELECT ON msdb.dbo.sysjobs TO gguser;			
User Database Roles and Privileges						
SYSADMIN	Yes	No	Required for a one time change to enable database level Change Data Capture (CDC) if not already enabled and can be revoked once TRANDATA has been enabled.			
			ALTER SERVER ROLE sysadmin ADD MEMBER gguser;			
			Database Administrators with sysadmin credentials can manually enable the database for CDC using the following, which would negate the need for the Extract user to have this privilege:			
			EXEC msdb.sys.sp_cdc_enable_db 'source_database'			
DBOWNER	Yes	Yes	ALTER ROLE db_owner ADD MEMBER gguser;			

Amazon RDS User Permissions and Requirements

The following user requirements and minimum database privileges and permissions are required for Oracle GoldenGate to capture from and apply to an Amazon RDS for SQL Server database:

- Create a dedicated login for Oracle GoldenGate for Amazon RDS for SQL Server.
- Add the login as a user to the msdb database and to the source or target database.
- 3. Create a schema in the source or target database, to be used for objects required for Oracle GoldenGate. This schema should map to the GGSCHEMA value used in the GLOBALS parameter file.
- **4.** Enable the following privileges and permissions for the Oracle GoldenGate user based on whether the user is for an Extract, or for a Replicat.

Table 21-2 Privileges and Permissions for Oracle GoldenGate User

Privilege	Extract	Replicat	Syntax		
msdb Database Roles and Privileges					
EXECUTE ON rds_cdc_enable_db	Yes	No	GRANT EXECUTE ON msdb.dbo.rds_cdc_enable_db TO gguser; Database administrators with master credentials can manually enable the database for Change Data Capture using the following command, which would negate the need for the Extract user to have this permission: EXEC msdb.dbo.rds_cdc_enable_db 'source_database'		
SQLAgentOperatorR ole	Yes	No	ALTER ROLE SQLAgentOperatorRole ADD MEMBER gguser;		
SQLAgentUserRole	Inherited	Yes	ALTER ROLE SQLAgentUserRole ADD MEMBER gguser;		
SELECT ON sysjobactivity	Yes	No	GRANT SELECT ON msdb.dbo.sysjobactivity TO gguser;		
SELECT ON sysjobs	Yes	No	GRANT SELECT ON msdb.dbo.sysjobs TO gguser;		
User Database Roles and Privileges					
DBOWNER	Yes	Yes	ALTER ROLE db_owner ADD MEMBER gguser;		

Azure SQL Database

The following user requirements and minimum database privileges and permissions are required for Oracle GoldenGate to apply to an Azure SQL Database:

- 1. Create a dedicated login for Oracle GoldenGate for Azure SQL Database.
- 2. Add the login as a user to the target database.
- Create a schema in the target database, to be used for objects required for Oracle GoldenGate. This schema should map to the GGSCHEMA value used in the GLOBALS parameter file.
- 4. Enable the following privileges and permissions for the Oracle GoldenGate user.

Table 21-3 Privileges and Permissions for Oracle GoldenGate User

Privilege	Extract	Replicat	Syntax
User Database Roles a	nd Privileges		
DBOWNER	NA	Yes	ALTER ROLE db_owner ADD MEMBER gguser;



Configuring a Database Connection

This section contains instructions for setting up the Extract and Replicat connections to a SQL Server database.

- Configuring an Extract Database Connection
- Configuring a Replicat Database Connection
- Configuring a Database Connection on Linux
- Configuring a Database Connection on Windows

Configuring an Extract Database Connection

Extract connects to a source SQL Server database through an ODBC (Open Database Connectivity) connection. To create this connection, set up a data source name (DSN) through the Data Sources (ODBC) control panel. See Configuring a Database Connection on Windows for instructions.

Connecting to the Listener in an Always On Environment

Extract can connect to the Listener of an Always On environment and can be configured to route its read-only queries to an available readable, synchronous mode secondary replica. By default, if Extract connects to a Listener, all processing will be done against the primary replica, but if the Extract is configured with the TRANLOGOPTIONS ALWAYSONREADONLYROUTING parameter, its read-only queries are routed by the Listener to an available readable secondary replica. See TRANLOGOPTIONS and Requirements Summary for Capture and Delivery of Databases in an AlwaysOn Availability Group for more information.

Configuring a Replicat Database Connection

Replicat can connect to the target database to perform DML operations in the following ways:

- Through ODBC.
- Through OLE DB if the SQL Server driver supports it.
- Through OLE DB as the SQL Server replication user. NOT FOR REPLICATION must be set on IDENTITY columns, foreign key constraints, and triggers.

Before you select a method to use, review the following guidelines and procedures to evaluate the advantages and disadvantages of each.

- Using ODBC or Default OLE DB
- Using OLE DB with USEREPLICATIONUSER

Using ODBC or Default OLE DB

If Replicat connects through ODBC or through the default OLE DB connection, the following limitations apply:

- To keep IDENTITY columns identical on source and target when using ODBC or default
 OLE DB, Replicat creates special operations in its transaction to ensure that the seeds are
 incremented on the target. These steps may reduce delivery performance.
- You must adjust or disable triggers and constraints on the target tables to eliminate the potential for redundant operations.



To use Replicat with either ODBC or OLE DB, follow these steps:

- To use ODBC exclusively, include the DBOPTIONS parameter with the USEODBC option in the Replicat parameter file. (To use the default OLE DB connection, no parameter is required.). For SQL Server CDC for Linux for Oracle GoldenGate, the USEODBC option is not allowed.
- 2. Disable triggers and constraints on the target tables. See Disabling Triggers and Cascade Constraints on the Target.
- 3. To use IDENTITY columns in a bidirectional SQL Server configuration, define the IDENTITY columns to have an increment value equal to the number of servers in the configuration, with a different seed value for each one. For example, a two-server installation would be as follows:
 - Sys1 sets the seed value at 1 with an increment of 2.
 - Sys2 sets the seed value at 2 with an increment of 2.

A three-server installation would be as follows:

- Sys1 sets the seed value at 1 with an increment of 3.
- Sys2 sets the seed value at 2 with an increment of 3.
- Sys3 sets the seed value at 3 with an increment of 3.
- 4. Configure an ODBC data source. See Configuring a Database Connection on Windows.

Using OLE DB with USEREPLICATIONUSER

If Replicat connects as the SQL Server replication user through OLE DB with the USEREPLICATIONUSER option, and NOT FOR REPLICATION is enabled for IDENTITY, triggers with foreign key constraints, the following benefits and limitations apply.

- IDENTITY seeds are not incremented when Replicat performs an insert. For SQL Server bidirectional configurations, stagger the seed and increment values like the example in Step 3 of the previous section.
- Triggers are disabled for the Replicat user automatically on the target to prevent redundant operations. However triggers fire on the target for other users.
- Foreign key constraints are not enforced on the target for Replicat transactions. CASCADE updates and deletes are not performed. These, too, prevent redundant operations.
- CHECK constraints are not enforced on the target for Replicat transactions. Even though
 these constraints are enforced on the source before data is captured, consider whether
 their absence on the target could cause data integrity issues.



Normal IDENTITY, trigger, and constraint functionality remains in effect for any users other than the Replicat replication user.

To use Replicat with USEREPLICATIONUSER, follow these steps:



Note:

For SQL Server CDC for Linux for Oracle GoldenGate, the USEREPLICATIONUSER option is not allowed.

Note:

For Replicat, connections using a Microsoft ODBC driver, install the Microsoft OLE DB Driver 18 for SQL Server to support the USEREPLICATIONUSER option: https://www.microsoft.com/en-us/download/details.aspx?id=56730

- 1. In SQL Server Management Studio (or other interface) set the NOT FOR REPLICATION flag on the following objects. For active-passive configurations, set it only on the passive database. For active-active configurations, set it on both databases.
 - Foreign key constraints
 - Check constraints
 - IDENTITY columns
 - Triggers (requires textual changes to the definition; see the SQL Server documentation for more information.)
- 2. Partition IDENTITY values for bidirectional configurations.
- 3. In the Replicat MAP statements, map the source tables to appropriate targets, and map the child tables that the source tables reference with triggers or foreign-key cascade constraints. Triggered and cascaded child operations are replicated by Oracle GoldenGate, so the referenced tables must be mapped to appropriate targets to preserve data integrity. Include the same parent and child source tables in the Extract TABLE parameters.

Note:

If referenced tables are omitted from the MAP statements, no errors alert you to integrity violations, such as if a row gets inserted into a table that contains a foreign key to a non-replicated table.

- 4. In the Replicat parameter file, include the DBOPTIONS parameter with the USEREPLICATIONUSER option.
- 5. Configure an ODBC data source. See Configuring a Database Connection on Windows.

Configuring a Database Connection on Linux

Oracle GoldenGate for SQL Server database configuration provides the same support for Linux and Windows. However, you need the <code>msodbcsql13*</code> or <code>msodbcsql17*</code> drivers to set up the connections in a Linux environment.

See:



https://docs.microsoft.com/en-us/sql/connect/odbc/linux-mac/installing-the-microsoft-odbc-driver-for-sql-server?view=sql-server-2017

https://docs.microsoft.com/en-us/sql/connect/odbc/download-odbc-driver-for-sql-server?view=sql-server-2017

The following example demonstrates how to create an ODBC data source in a Linux environment:

1. Create a template file for your data source:

```
vi odbc template file.ini
```

2. Describe the data source in the template file. In the following example, myserver_ss2017_source is used as the name with DBLOGIN and SOURCEDB to connect to the database.:

```
[myserver_SS2017_source]
Driver = ODBC Driver 17 for SQL Server
Server = myserver,1433
Database = source_database
User = ssuser
Password = ssuserpassword
```

3. Install the data source using the command:

```
odbcinst -i -s -f odbc template file.ini
```

For more information, see:

https://docs.microsoft.com/en-us/sql/connect/odbc/linux-mac/connection-string-keywords-and-data-source-names-dsns?view=sql-server-2017

Configuring a Database Connection on Windows

Follow these instructions to create a SQL Server system data source name (DSN) for a source or a target SQL Server database. A DSN stores information about how to connect to a SQL Server database through ODBC (Open Database Connectivity).

To create a SQL Server DSN

- To run the ODBC client, select Control Panel, select Administrative Tools, and then select Data Sources (ODBC).
- In the ODBC Data Source Administrator dialog box of the ODBC client, select the System DSN tab, and then click Add.
- Under Create New Data Source, select the correct SQL Server driver supported for your version of SQL Server, and then click Finish. The Create a New Data Source to SQL Server wizard appears. To determine the correct SQL Server driver, see Database Connection.
- Supply the following, and then click Next:
 - Name: Can be of your choosing. In a Windows cluster, use one name across all nodes in the cluster.
 - Description: (Optional) Type a description of this data source.
 - Server: Select the SQL Server server or instance name. Optionally, the listener\instance name of an Always On Availability Group can be listed.
- 5. For login authentication, do one of the following, and then click **Next**:



- Select With Integrated Windows Authentication for Oracle GoldenGate to use Windows authentication.
- b. To use database credentials, With SQL Server authentication using a login ID and password entered by the user, and supply login information.
- 6. If the default database is not set to the one that Oracle GoldenGate will connect to, select **Change the default database to**, and then select the database. Set the other settings to use ANSI. Click **Next**.
- 7. Leave the next page set to the defaults. Click Finish.
- 8. Click **Test Data Source** to test the connection.
- 9. If the test is successful, close the confirmation box and the Create a New Data Source box.
- 10. Repeat this procedure for each SQL Server source and target system.

Preparing Tables for Processing

The table attributes in the following sections must be addressed in your Oracle GoldenGate environment.

- Disabling Triggers and Cascade Constraints on the Target
- · Assigning Row Identifiers
- Improving IDENTITY Replication with Array Processing

Disabling Triggers and Cascade Constraints on the Target

In an environment where SQL Server is the target, consider triggers and cascade constraints that may repeat an operation that occurred on the source. For example, if the source has an insert trigger on TableA that inserts a record into TableB, and Oracle GoldenGate is configured to capture and deliver both TableA and TableB, the insert trigger on the target table, TableA, must be disabled. Otherwise, Replicat inserts into TableA, and the trigger fires and insert into TableB. Replicat will then try to insert into TableB, and then terminate abnormally.

When a trigger or cascade constraint repeats an operation that occurred on the source, you do not have to disable the trigger or constraint when the following conditions are both true:

- You use the DBOPTIONS USEREPLICATIONUSER parameter in Replicat.
- You use OLE DB connection for Replicat. The use of the OLE DB connection is the default configuration. Note that the trigger, constraint, or IDENTITY property must have NOT FOR REPLICATION enabled.

In the following scenario, disable the triggers and constraints on the target:

Uni-directional replication where all tables on the source are replicated.

In the following scenarios, enable the triggers and constraints on the target:

- Uni-directional replication where tables affected by a trigger or cascade operation are not replicated, and the only application that loads these tables is using a trigger or cascade operation.
- Uni-directional or -bi-directional replication where all tables on the source are replicated. In this scenario, set the target table cascade constraints and triggers to enable NOT FOR REPLICATION, and use the DBOPTIONS USEREPLICATIONUSER parameter in Replicat.



Assigning Row Identifiers

Oracle GoldenGate requires unique row identifiers on the source and target tables to locate the correct target rows for replicated updates and deletes. Source tables can have any kind of key listed in How Oracle GoldenGate Determines the Kind of Row Identifier to Use. If there is no primary key identified on a table that has fixed-length columns, the length of one of the fixed-length columns must be below 3800 bytes.

- How Oracle GoldenGate Determines the Kind of Row Identifier to Use
- Using KEYCOLS to Specify a Custom Key

How Oracle GoldenGate Determines the Kind of Row Identifier to Use

Unless a KEYCOLS clause is used in the TABLE or MAP statement, Oracle GoldenGate selects a row identifier to use in the following order of priority:

- 1. First unique key alphanumerically that does not contain a timestamp or non-materialized computed column.
- 2. If neither of these key types exist, Oracle GoldenGate constructs a pseudokey of all columns that the database allows to be used in a unique key, excluding those that are not supported by Oracle GoldenGate in a key or those that are excluded from the Oracle GoldenGate configuration. For SQL Server, Oracle GoldenGate requires the row data in target tables that do not have a primary key to be less than 8000 bytes.



If there are types of keys on a table or if there are no keys at all on a table, Oracle GoldenGate logs a message to the report file. Constructing a key from all of the columns impedes the performance of Oracle GoldenGate on the source system. On the target, this key causes Replicat to use a larger, less efficient WHERE clause.

Using **KEYCOLS** to Specify a Custom Key

If a table does not have an applicable row identifier, or if you prefer that identifiers are not used, you can define a substitute key, providing that the table has columns that always contain unique values. You define this substitute key by including a KEYCOLS clause within the Extract TABLE parameter and the Replicat MAP parameter. The specified key overrides any existing primary or unique key that Oracle GoldenGate finds.

Improving IDENTITY Replication with Array Processing

Because only one table per session can have IDENTITY_INSERT set to ON, Replicat must continuously toggle IDENTITY_INSERT when it applies IDENTITY data to multiple tables in a session. To improve the performance of Replicat in this situation, use the BATCHSQL parameter. BATCHSQL causes Replicat to use array processing instead of applying SQL statements one at a time.



Globalization Support

Oracle GoldenGate provides globalization support that lets it process data in its native language encoding. The Oracle GoldenGate apply process (Replicat) can convert data from one character set to another when the data is contained in character column types.



Preparing the Database for Oracle GoldenGate — CDC Capture

Learn how to enable supplemental logging in the source database tables that are to be used for capture by the Extract for SQL Server and how to purge older CDC staging data.

You can learn more about CDC Capture with this Oracle By Example:

Using the Oracle GoldenGate for SQL Server CDC Capture Replication http://www.oracle.com/webfolder/technetwork/tutorials/obe/fmw/goldengate/12c/sql_cdcrep/sql_cdcrep.html.

Topics:

- · Enabling CDC Supplemental Logging
- Purging CDC Staging Data
- Enabling Bi-Directional Loop Detection

Enabling CDC Supplemental Logging

With the CDC Extract, the method of capturing change data is via SQL Server Change Data Capture tables, so it is imperative that you follow the procedures and requirements below, so that change data is correctly logged, maintained, and captured by Extract.

You will enable supplemental logging with the ADD TRANDATA command so that Extract can capture the information that is required to reconstruct transactions.

ADD TRANDATA must be issued for all tables that are to be captured by Oracle GoldenGate, and to do so requires that a valid schema be used in order to create the necessary Oracle GoldenGate tables and stored procedures.

Enabling supplemental logging for a CDC Extract does the following:

- Enables SQL Server Change Data Capture at the database level, if it's not already enabled.
 - EXECUTE sys.sp_cdc_enable_db
- Creates a Change Data Capture staging table for each base table enabled with supplemental logging by running EXECUTE sys.sp_cdc_enable_table, and creates a trigger for each CDC table. The CDC table exists as part of the system tables within the database and has a naming convention like, cdc.OracleGG basetableobjectid CT.
- Creates a tracking table of naming convention schema.OracleGGTranTables. This table is
 used to store transaction indicators for the CDC tables, and is populated when the trigger
 for a CDC table is fired. The table will be owned by the schema listed in the GLOBALS file's,
 GGSCHEMA parameter.
- Creates a unique fetch stored procedure for each CDC table, as well as several other stored procedures that are required for Extract to function. These stored procedures will be owned by the schema listed in the GLOBALS file's, GGSCHEMA parameter.
- Also, as part of enabling CDC for tables, SQL Server creates two jobs per database:

```
cdc.dbname_capture
cdc.dbname cleanup
```

- The CDC Capture job is the job that reads the SQL Server transaction log and populates
 the data into the CDC tables, and it is from those CDC tables that the Extract will capture
 the transactions. So it is of extreme importance that the CDC capture job be running at all
 times. This too requires that SQL Server Agent be set to run at all times and enabled to run
 automatically when SQL Server starts.
- Important tuning information of the CDC Capture job can be found in CDC Capture Method Operational Considerations.
- The CDC Cleanup job that is created by Microsoft does not have any dependencies on whether the Oracle GoldenGate Extract has captured data in the CDC tables or not.
 Therefore, extra steps need to be followed in order to disable or delete the CDC cleanup job immediately after TRANDATA is enabled, and to enable Oracle GoldenGate's own CDC cleanup job. See Retaining the CDC Table History Data for more information.

The following steps require a database user who is a member of the SQL Server System Administrators (sysadmin) role.

- 1. In the source Oracle GoldenGate installation, ensure that a GLOBALS (all CAPS and no extension) file has been created with the parameter GGSCHEMA schemaname. Ensure that the schema name used has been created (CREATE SCHEMA schemaname) in the source database. This schema will be used by all subsequent Oracle GoldenGate components created in the database, therefore it is recommended to create a unique schema that is solely used by Oracle GoldenGate, such as 'ogg'. It is recommended not to use the SQL Server schema cdc and to create a new schema specific to Oracle GoldenGate.
- 2. On the source system, run GGSCI
- 3. Issue the following command to log into the database:

DBLOGIN SOURCEDB DSN [, {USERID user, PASSWORD password | USERIDALIAS alias}] Where:

- SOURCEDB DSN is the name of the SQL Server data source.
- USERID user is the database login and PASSWORD password is the password that is required if the data source connects via SQL Server authentication.

 Alternatively, USERIDALIAS alias is the alias for the credentials if they are stored in a credentials store. If using DBLOGIN with a DSN that is using Integrated Windows authentication, the connection to the database for the GGSCI session will be that of the user running GGSCI. In order to issue ADD TRANDATA or DELETE TRANDATA, this user must be a member of the SQL Server sysadmin server role.
- 4. In GGSCI, issue the following command for each table that is, or will be, in the Extract configuration. You can use a wildcard to specify multiple table names.

```
ADD TRANDATA owner.*
```

Optionally, you can designate the filegroup in which the SQL Server Change Data Capture staging tables will be placed, by using the FILEGROUP option with an existing filegroup name.

```
ADD TRANDATA owner.table FILEGROUP cdctables See ADD TRANDATA
```



Purging CDC Staging Data

When enabling supplemental logging, data that is required by Extract to reconstruct transactions are stored in a series of SQL Server CDC system tables, as well Oracle GoldenGate objects that are used to track operations within a transaction. And as part of enabling supplemental logging, SQL Server will create its own Change Data Capture Cleanup job that runs nightly by default, and purges data older than 72 hours. The SQL Server CDC Cleanup job is unaware that an Extract may still require data from these CDC system tables and can remove that data before the Extract has a chance to capture it.

If data that Extract needs during processing has been deleted from the CDC system tables, then one of the following corrective actions might be required:

- Alter Extract to capture from a later point in time for which CDC data is available (and accept possible data loss on the target).
- Resynchronize the source and target tables, and then start the Oracle GoldenGate environment over again.

To remedy this situation, Oracle GoldenGate for SQL Server includes Oracle GoldenGate for SQL Server includes the ogg_cdc_cleanup_setup.bat program that is used to create an Oracle GoldenGate Cleanup job associated stored procedures and tables.

The Extract, upon startup, will expect, by default, that those Oracle GoldenGate Cleanup task objects exist and will stop if they do not. Extract will issue a warning if the SQL Server CDC Cleanup job exists alongside the Oracle GoldenGate Cleanup job.

The default checks by Extract for the Oracle GoldenGate CDC Cleanup task objects can be overwritten by using the TRANLOGOPTIONS NOMANAGECDCCLEANUP in the Extract, but this would only be recommended for development and testing purposes.

Use the following steps immediately after enabling supplemental logging and prior to starting the Extract, to create the Oracle GoldenGate CDC Cleanup job and associated objects. You can re-run these steps to re-enable this feature should any of the objects get manually deleted.

To create the Oracle GoldenGate CDC Cleanup job and objects:

The ogg cdc cleanup setup file is located in the the home directory for Classic Architecture.

The script uses the Microsoft sqlemd utility, so ensure that sqlemd is installed on the system where Oracle GoldenGate is installed.

This requires an SQL Server authenticated database user who is a member of the SQL Server System Administrators (sysadmin) role. Windows authentication is not supported for the .bat script.

Stop and disable the database's SQL Server cdc. dbname_cleanup job from SQL Server
Agent. Alternatively, you can drop it from the source database with the following command.

```
EXECUTE sys.sp cdc drop job 'cleanup'
```

2. Run the ogg cdc cleanup setup.bat file, providing the following variable values.

For Windows:

ogg_cdc_cleanup_setup.bat createJob userid password databasename
servername\instancename schema



For Linux:

```
./ogg_cdc_cleanup_setup.sh createJob userid password databasename
"servername,port" schema
```

In the preceding examples, USER ID and password should be a valid SQL Server login and password for a user, which has sysadmin rights. The databasename, servername\instancename, or servername, port, are the source database name, server,

servername\instancename, or servername, port, are the source database name, server, and instance, or server and TCP/IP port where SQL Server is running. If only the server name is listed, then the default instance will be connected to. The schema is the schema name listed in the GLOBALS file, with the GGSCHEMA parameter. This schema should be the same for all Oracle GoldenGate objects, including supplemental logging, checkpoint tables, heartbeat tables, and the Oracle GoldenGate CDC Cleanup job.

For example:

```
ogg_cdc_cleanup_setup.bat createJob ggsuser ggspword db1 server1\inst1 ogg
```

When using server, port in the connection string, enclose the string in double quotes, for example:

```
ogg_cdc_cleanup_setup.bat createJob login password source database
"sql2016.samplestring.us-west-1.rds.amazonaws.com,1433" OGG schema name
```

The Oracle GoldenGate CDC Cleanup job when created, is scheduled to run every ten minutes, with a default retention period of seventy two hours. The job will not purge data for an Extract's recovery checkpoint however, regardless of the retention period.

Additional information of the Oracle GoldenGate CDC Cleanup job can be found in CDC Capture Method Operational Considerations.

Enabling Bi-Directional Loop Detection

Loop detection is a requirement for bi-directional implementations of Oracle GoldenGate, so that an Extract for one source database does not recapture transactions sent by a Replicat from another source database.

With the CDC Extract capture method, by default, any transaction committed by a Replicat into a database where an Extract is configured, will recapture that transaction from the Replicat as long as supplemental logging is enabled for those tables that the Replicat is delivering to.

In order to ignore recapturing transactions that are applied by a Replicat, you must use the TRANLOGOPTIONS FILTERTABLE parameter for the CDC Extract. The table used as the filtering table will be the Oracle GoldenGate checkpoint table that you must create for the Replicat.

To create a Filter Table and enable Supplemental Logging:

The steps below require a database user who is a member of the SQL Server System Administrators (sysadmin) role.

- 1. On the source system, run GGSCI
- 2. Issue the following command to log into the database.

```
DBLOGIN SOURCEDB DSN [,{USERID user, PASSWORD password | USERIDALIAS alias}]
```



In the preceding example, the <code>SOURCEDB DSN</code> is the name of the SQL Server data source. The <code>USERID</code> user is the database login and <code>PASSWORD</code> password is the password that is required if the data source connects through SQL Server authentication.

Alternatively, <code>USERIDALIAS</code> alias is the alias for the credentials if they are stored in a credentials store. If using <code>DBLOGIN</code> with a DSN that is using Integrated Windows authentication, the connection to the database for the GGSCI session is that of the user running GGSCI. In order to issue <code>ADD TRANDATA</code> or <code>DELETE TRANDATA</code>, this user must be a member of the SQL Server <code>sysadmin</code> server role.

3. Create the Oracle GoldenGate checkpoint table that is used by the Replicat to deliver data to the source database.

Example: ADD CHECKPOINTTABLE ogg.ggchkpt

It is recommended that you use the same schema name as used in the GGSCHEMA parameter of the GLOBALS file.

4. Enable supplemental logging for the newly created checkpoint table.

Example: ADD TRANDATA ogg.ggchkpt

5. Add the Replicat with the checkpoint table information.

Example: ADD REPLICAT reptgt1, EXTTRAIL ./dirdat/e2,checkpointtable ogg.ggchkpt

6. Configure the Extract with the IGNOREREPLICATES (on by default) and FILTERTABLE parameters, using the Replicat's checkpoint table for the filtering table.

TRANLOGOPTIONS IGNOREREPLICATES

TRANLOGOPTIONS FILTERTABLE ogg.ggchkpt



Requirements Summary for Capture and Delivery of Databases in an Always On Availability Group

Oracle GoldenGate for SQL Server supports capture from a primary replica or a read-only, synchronous mode secondary replica of an Always On Availability Group, and delivery to the primary replica.

When capturing from either a primary or a secondary replica in an Always On Availability Group, it is important to understand that the capture process must only read hardened transactions from the log, and that there be no potential for data loss between any replica database that Oracle GoldenGate is or will capture from.

Topics:

- Database Connection
- Supplemental Logging
- Operational Requirements and Considerations

Database Connection

For both Extract and Replicat, it is recommended to create a System DSN that uses the Always On Availability Group Listener for the connection.

- For the Replicat, connecting to the Listener allows the Replicat to reconnect if the primary replica performs a failover to a new instance, without having to manually edit the DSN settings to point to the new primary.
- For the Extract connecting to the Listener not only allows reconnecting to the primary
 without editing the DSN to point to the new instance, but also provides the optional ability
 to run the Extract's data extraction stored procedures, against a read-only secondary.
- For both Extract and Replicat connected to an Always On environment, use the AUTORESTART parameter for the Manager, to restart the processes after a failover.
- To route the Extract's data extraction queries to a read-only secondary, ensure that the DSN connection uses the Listener, that you have one or more read-only secondary replicas that are configured to handle read-only routing, and that the Extract runs with the TRANLOGOPTIONS ALWAYSONREADONLYROUTING parameter.
 - Ensure that the Application Intent field of the DSN configuration is set to READWRITE and not READONLY
 - Refer to the following Microsoft documentation on how to configure read-only routing: https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/configure-read-only-routing-for-an-availability-group-sql-server?view=sql-server-2017

Supplemental Logging

Supplemental logging must be enabled by normal means (ADD TRANDATA) using GGSCI connected to the primary replica and not against a secondary replica.

- Create a DSN to the primary replica, or to the Always On Availability Group Listener, to connect using DBLOGIN to run ADD TRANDATA from GGSCI.
- The login used to enable supplemental logging must have sysadmin membership of the primary replica instance.
- When enabling supplemental logging against the primary replica database, the SQL Server Change Data Capture job does not automatically get created on any secondary replicas.
 Upon failover from a primary to a secondary, you must manually create the SQL Server Change Data Capture job and the Oracle CDC Cleanup job if in use, on the new primary replica.

```
EXECUTE sys.sp cdc add job N'capture
```

When creating the SQL Server CDC Capture job on the new primary, the default configuration settings are put in place. So if you have previously modified the default values on the former primary replica, you need to run sys.sp_cdc_change_job on the new primary and set the values accordingly.

Note:

Consult the Microsoft documentation on how to enable the CDC Capture job for AlwaysOn Secondary Replicas for more information.

Operational Requirements and Considerations

- When an instance is no longer the primary instance but has the SQL Server CDC Capture
 job installed, the job ceases to run after some time and does not attempt to restart. Upon
 failover back to that instance, the job does not automatically start, so it must be manually
 started.
- If secondary replica databases are not in sync with the primary replica database, the CDC capture job will not advance in the log, and therefore no records will be captured by an Extract, until such time that the primary and secondary replicas are synchronized. See this article from Microsoft for more details:

https://docs.microsoft.com/en-us/sql/database-engine/availability-groups/windows/replicate-track-change-data-capture-always-on-availability?view=sql-server-2017

Note:

When capturing from either a primary or a secondary replica in an Always On Availability Group, it is important to understand that the capture process must only read hardened transactions from the log, and that there be no potential for data loss between any replica database that Oracle GoldenGate is or will capture from.



- When running an Extract from a middle tier Windows or Linux server, set the middle tier server's date, time, and time zone to the same as that of the primary replica.
- Upon failover from a primary to a secondary replica, reinstall the Oracle GoldenGate CDC Cleanup job on the new primary by re-running the <code>ogg_cdc_cleanup_setup.bat</code> file with the <code>createJob</code> option.
- If Extract is configured to capture from a readable secondary replica, but not configured with read-only routing, the SQL Server CDC Capture job must be created against the secondary replica prior to starting the Extract, as the Extract will check if the job exists. To create the SQL Server CDC Capture job, any potential secondary that will have an Extract connected to it, must at some point be set to a writable Primary database and then follow the steps above, under supplemental logging, to manually add the SQL Server CDC Capture job.
- If uninstalling Oracle GoldenGate and disabling Change Data Capture on a database that
 is part of an Always On availability group, follow the extra steps provided in Disabling
 Change Data Capture.



CDC Capture Method Operational Considerations

This section provides information about the SQL Server CDC Capture options, features, and recommended settings.

Topics:

- Tuning SQL Server Change Data Capture
- Oracle GoldenGate CDC Object Versioning
- Valid and Invalid Extract Parameters for SQL Server Change Data Capture
- Details of the Oracle GoldenGate CDC Cleanup Process
- Changing from Classic Extract to a CDC Extract
- Restoring a Source Database Keeping CDC Data

Tuning SQL Server Change Data Capture

The following information is useful in improving the capture performance of the Extract.

- Ensure that Auto Create Statistics and Auto Update Statistics are enabled for the
 database. Maintaining statistics on the cdc.OracleGG_######_CT , cdc.lsn_time_mapping,
 and OracleGGTranTables table is crucial to the performance and latency of the Extract.
- The SQL Server Change Data Capture job collects data from the SQL Server transaction log and loads it into the Change Data Capture staging tables within the database.

As part of the job that is created, there are several available tuning parameters that can be used, and information on how to best tune the job can be found in the following article: https://technet.microsoft.com/en-us/library/dd266396(v=sql.100).aspx

As a general recommendation, you should change the SQL Server Change Data Capture Job polling interval from the default of 5 seconds to 1 second.

To change the default polling interval of the CDC Capture job, execute the following queries against the database:

```
EXEC [sys].[sp_cdc_change_job]
@job_type = N'capture',
@pollinginterval = 1,
GO,
--stops cdc job
EXEC [sys].[sp_cdc_stop_job],
@job_type = N'capture',
GO,
--restarts cdc job for new polling interval to take affect
EXEC [sys].[sp_cdc_start_job],
@job_type = N'capture',
```

Oracle GoldenGate CDC Object Versioning

Oracle GoldenGate provides a version tracking subsystem to track the CDC objects that are created by Oracle GoldenGate when enabling supplemental logging. These objects are:

- Oracle GoldenGate change tracking tables in the format OracleGG object id CT.
- Stored procedures in the format fetch_database name_object id
- Stored procedures OracleCDCExtract, OracleGGCreateProcs, and OracleGGCreateNextBatch.
- After successfully completing the ADD TRANDATA command, GGSCI creates a table called OracleGGVersion under the GGSCHEMA specified in the GLOBALS file, if it does not already exist.

Next, GGSCI inserts a record into the table that tracks the start and end time of the TRANDATA session. When Extract starts up, it checks for consistency between itself and the Oracle GoldenGate CDC objects by comparing its internal version number with the version numbers found in the OracleGGVersion table. If it finds that the version numbers do not match, it abends with a message similar to the following:

ERROR OGG-05337 The Oracle GoldenGate CDC object versions on database, source, are not consistent with the expected version, 2. The following versions(s) were found: 1. Rerun ADD TRANDATA for all tables previously enabled, including heartbeat, heartbeat seed, and filter tables.

Valid and Invalid Extract Parameters for SQL Server Change Data Capture

This section describes parameters used for the CDC Capture method. For more information about supported and unsupported parameters for the CDC Capture method, review *Reference for Oracle GoldenGate*.

TRANLOGOPTIONS LOB CHUNK SIZE

The Extract parameter LOB_CHUNK_SIZE is added for the CDC Capture method to support large objects. If you have huge LOB data sizes, then you can adjust the LOB_CHUNK_SIZE from the default of 4000 bytes, to a higher value up to 65535 bytes, so that the fetch size is increased, reducing the trips needed to fetch the entire LOB.

Example: TRANLOGOPTIONS LOB_CHUNK_SIZE 8000

TRANLOGOPTIONS MANAGECDCCLEANUP/NOMANAGECDCCLEANUP

The Extract parameter MANAGECDCCLEANUP/NOMANAGECDCCLEANUP is used by the CDC Capture method to instruct the Extract on whether or not to maintain recovery checkpoint data in the Oracle GoldenGate CDC Cleanup job. The default value is MANAGECDCCLEANUP and it doesn't have to be explicitly listed in the Extract. However, it does require creating the Oracle GoldenGate CDC Cleanup job prior to starting the Extract. MANAGECDCCLEANUP should be used for all production environments, where NOMANAGECDCCLEANUP may be used for temporary and testing implementations as needed.

Example: TRANLOGOPTIONS MANAGECDCCLEANUP

TRANLOGOPTIONS EXCLUDEUSER/EXCLUDETRANS

The SQL Server CDC Capture job does not capture user information or transaction names associated with a transaction, and as this information is not logged in the CDC staging tables, Extract has no method of excluding DML from a specific user or DML of a specific transaction name. The EXCLUDEUSER and EXCLUDETRANS parameters are therefore not valid for the CDC Capture process.

TRANLOGOPTIONS MANAGESECONDARYTRUNCATIONPOINT/NOMANAGESECONDARYTRUNCATIONPOINT/ACTIVESECONDARYTRUNCATIONPOINT

The SQL Server Change Data Capture job is the only process that captures data from the transaction log when using the Oracle GoldenGate CDC Capture method. The secondary truncation point management is not handled by the Extract, and for the Change Data Capture Extract, these parameters are not valid.

TRANLOGOPTIONS ALWAYSONREADONLYROUTING

The ALWAYSONREADONLYROUTING parameter allows Extract for SQL Server to route its read-only processing to an available read-intent Secondary when connected to an Always On availability group listener.

TRANLOGOPTIONS QUERYTIMEOUT

Specifies how long queries to SQL Server will wait for results before reporting a timeout error message. This option takes an integer value to represent the number of seconds. The default query timeout value is 300 seconds (5 minutes). The minimum value is 0 seconds (infinite timeout). The maximum is 2147483645 seconds.

TRANLOGOPTIONS TRANCOUNT

Allows adjustment of the number of transactions processed per each call by Extract to pull data from the SQL Server change data capture staging tables. Based on your transaction workload, adjusting this value may improve capture rate throughput, although not all workloads will be positively impacted. The minimum value is 1, maximum is 100, and the default is 10.

Details of the Oracle GoldenGate CDC Cleanup Process

The Oracle GoldenGate CDC Cleanup job is required for a CDC Extract by default, since Extract defaults to TRANLOGOPTIONS MANAGECDCCLEANUP. It is installed from a Windows batch file (ogg_cdc_cleanup_setup.bat), which uses sqlcmd to connect to the source SQL Server database and create the necessary objects and job.

There should be one job for each database enabled for CDC Capture, and you must create the job and objects following the steps mentioned in the Preparing the Database for Oracle GoldenGate — CDC Capture section of this document.

Additional options for the utility are discussed in the following sections.

The steps below require a SQL Server authenticated database user who is a member of the SQL Server System Administrators (sysadmin) role. Windows authentication is not supported for the .bat batch file.



Removing an Extract from the Database

When the Oracle GoldenGate CDC Cleanup object tables exist, each CDC Extract that is started against that database will create an entry in the <code>OracleGGExtractCheckpoint</code> table. This entry tracks a particular Extract's point in time recovery checkpoint, which is used as the cutoff LSN for the Oracle GoldenGate CDC cleanup tasks. If there are multiple Extracts running, each logging more recent recovery checkpoints in the table, but one Extract has been removed from the system without removing its entry into the <code>OracleGGExtractCheckpoint</code> table, then no data will be purged newer than that deleted Extract's old recovery checkpoint for all of the CDC staging tables. So when deleting an Extract from the database, follow the steps below to remove the Extract from the <code>OracleGGExtractCheckpoint</code> table if more than one Extract is running against the database.

- 1. Ensure that the Extract has been deleted using DELETE EXTRACT extract name
- **2.** On the source system, open a command prompt and change to the Oracle GoldenGate installation folder.
- 3. Run the ogg cdc cleanup setup.bat/.sh file, providing the following variable values:

ogg_cdc_cleanup_setup.bat deleteExtCheckpoint userid password databasename servername\instancename schema

Example: ogg_cdc_cleanup_setup.bat deleteExtCheckpoint ggsuser ggspword db1 server1\inst1 ogg

4. You will be prompted to enter the name of the Extract that is to be removed, and once entered, press the Enter/Return key to continue.

Modifying the Oracle GoldenGate CDC Cleanup Job

The default schedule, retention period and operation batch size for the Oracle GoldenGate CDC Cleanup job of a database is to run every 10 minutes, with a data retention policy of 72 hours (listed as 4320 minutes), purging in batches of 500 records per transaction until the retention policy is meet, not to exceed the recovery checkpoint data of the Extract.

For variations in customer environments, or change data table data retention requirements, it may be necessary to adjust these properties to increase the purge batch size or to adjust retention policies and the job run-time schedule.

To adjust the job execution frequency, manually modify the schedule for the <code>OracleGGCleanup_dbname_Job</code> job within SQL Server Agent. If you need to adjust the retention period or purge batch size, you must manually edit the job step for the <code>OracleGGCleanup_dbname_Job</code> job within SQL Server Agent. The job step passes two parameters to the cleanup stored procedure, and you can modify the value for <code>@retention_minutes</code> to adjust the data retention policy as needed, or modify the <code>@threshold</code> value to increase or decrease the <code>purge</code> batch size. In high transactional environments, it may be necessary to increase the <code>@threshold</code> value to a number such as 10000. Monitoring the amount of time that it takes for the job to run within each cycle can be used to determine effective <code>@threshold</code> values.

Deleting the Oracle GoldenGate CDC Cleanup Job

If you no longer require the Oracle GoldenGate CDC Cleanup job and associated objects and need to remove them, perform the following steps:

- 1. Open a command prompt and change to the Oracle GoldenGate installation folder.
- **2.** Run the ogg cdc cleanup setup.bat file, providing the following variable values:



ogg_cdc_cleanup_setup.bat dropJob userid password databasename
servername\instancename schema

 $\textbf{Example:} \verb| ogg_cdc_cleanup_setup.bat dropJob ggsuser ggspword db1 server1 | linst1 ogg the line of the line$

Changing from Classic Extract to a CDC Extract

If you plan to change from using a Classic Extract from Oracle GoldenGate 12c (12.3.0.1) or earlier, to an Oracle GoldenGate 19c CDC Extract, then you must remove the supplemental logging that was implemented using the Classic Extract installation method, and re-enable supplemental logging using the CDC Extract installation binaries, as the calls to enable TRANDATA are different between the two versions, and the implementation of TRANDATA for Classic Extract is not supported by the CDC Extract.

Follow these general guidelines to remove and re-enable supplemental logging. Special consideration and planning should be involved if migrating from Classic to CDC Extract in a production system. The information provided here does not cover all requirements and is only offered as general requirements regarding supplemental logging:

- 1. Ensure that the Classic Extract has processed all remaining data in the logs and can be gracefully stopped.
- Do one of the following, depending on how Extract was running in relation to other replication or CDC components:
 - If Extract was not running concurrently with SQL Server transactional replication or a non-Oracle CDC configuration on the same database, open a query session in Management Studio and issue the following statement against the source database to disable and delete any CDC or replication components, and to clear the secondary truncation point.

```
EXEC sys.sp_cdc_disable_db
```

 If Extract was running concurrently with SQL Server transactional replication or a non-Oracle CDC configuration on the same database, run GGSCI from the Classic Extract's installation folder, login to the source database with the DBLOGIN, and then issue the following command for each table that is in the Extract configuration. You can use a wildcard to specify multiple table names

```
DELETE TRANDATA owner.table
DELETE TRANDATA owner.*
```

3. Delete any heartbeat table entries if one was installed.

```
DELETE HEARTBEATTABLE
```

4. Using the Oracle GoldenGate CDC Extract installation binaries, follow the steps listed in Preparing the Database for Oracle GoldenGate — CDC Capture to re-enable supplemental logging and other necessary components, and re-add the heartbeat table.

Restoring a Source Database Keeping CDC Data

When restoring a SQL Server database that has been enabled with CDC and you want to keep the existing CDC staging data that has already accumulated in the database, as well as the CDC settings, you must specify the KEEP CDC option with the RESTORE statement.

This requirement is only if restoring the database to a new instance, or to the same instance but with a different database name. If you are restoring the original database on the instance, then the option is not required.

For further requirements and understanding, review the following document link from Microsoft:

https://docs.microsoft.com/en-us/sql/t-sql/statements/restore-statements-arguments-transact-sql?view=sql-server-ver15

Additionally, you must manually recreate the CDC Capture job and the Oracle GoldenGate CDC Cleanup job.



Part VIII

Using Oracle GoldenGate with Sybase

With Oracle GoldenGate for Sybase database, you can replicate data to and from supported Sybase versions or between a Sybase database and a database of another type. Oracle GoldenGate for Sybase supports data filtering, mapping, and transformation.

Use this part to get started with Oracle GoldenGate on a Sybase database system and performing initial setup. See Installing for Sybase in *Installing Oracle GoldenGate* for installation requirements.

Topics:

- Understanding What's Supported for Sybase
 This chapter contains information on database and table features supported by Oracle GoldenGate for Sybase.
- Preparing the System for Oracle GoldenGate



Understanding What's Supported for Sybase

This chapter contains information on database and table features supported by Oracle GoldenGate for Sybase.

Topics:

- Supported Sybase Data Types
- Non-Supported Sybase Data Types
- Supported Operations and Objects for Sybase
- Non-Supported Operations and Objects for Sybase

Supported Sybase Data Types

This section lists the Sybase data types that Oracle GoldenGate supports and any limitations of this support.

- Integers
- Floating-Point Numbers
- Character Data
- Dates and Timestamps
- Large Objects
- Money Types
- IDENTITY Type
- text, image, and unitext Data Types
- User-Defined Data Types

Integers

- BIGINT
- BIT
- DECIMAL
- INT (signed and unsigned)
- TINYINT (signed and unsigned)
- NUMERIC
- SMALLINT (signed and unsigned)

Limitations of Support

• NUMERIC and DECIMAL (fixed-point) are supported with no integrity loss when moving data to a target column of the same data type without involving calculations or transformation.

When calculations or transformation must be performed, Oracle GoldenGate supports a maximum value of a signed long integer (32-bits).

- BIT is supported for automatic mapping between Sybase databases. To move BIT data between Sybase and another database type, Oracle GoldenGate treats BIT data as binary. In this case, the following are required:
 - The BIT column must be mapped to the corresponding source or target column with a COLMAP clause in a TABLE or MAP statement.
- For the Sybase 157 GA release, these data types cannot be replicated:
 - BIGINT (as a key column)
 - BIGDATETIME
 - BIGTIME
- When replicating TINYINT and Extract is not in the same version of Replicat, you will need
 to create a sourcedef and/or targetdef file even if you are replicating between identical
 Sybase versions.
- See also Non-Supported Sybase Data Types.

Floating-Point Numbers

- DOUBLE
- FLOAT
- REAL

Limitations of Support

The support of range and precision for floating-point numbers depends on the host machine. In general, the precision is accurate to 16 significant digits, but you should review the database documentation to determine the expected approximations. Oracle GoldenGate rounds or truncates values that exceed the supported precision.

Character Data

- CHAR
- NCHAR
- NVARCHAR
- VARCHAR
- UNICHAR
- UNIVARCHAR

Limitations of Support

- These data types are supported to the maximum length supported by the database, this being the maximum page size.
- Fetching NVARCHAR replication results using the Sybase char_length or datalength functions when a Sybase database is the target and the source is a heterogenous database and you replicate from the source to the target may result in a data integrity issue. This occurs when you use a Sybase release earlier than Adaptive Server Enterprise 15.5 for Windows x64 platform EBF 21262: 15.5 ESD #5.3.



Dates and Timestamps

- BIGDATETIME
- BIGTIME
- DATE
- DATETIME
- SMALLDATETIME
- TIME

Limitations of Support

- Oracle GoldenGate supports timestamp data from 0001/01/03:00:00:00 to 9999/12/31:23:59:59. If a timestamp is converted from GMT to local time, these limits also apply to the resulting timestamp. Depending on the time zone, conversion may add or subtract hours, which can cause the timestamp to exceed the lower or upper supported limit.
- Oracle GoldenGate does not support negative dates.

Large Objects

- BINARY
- IMAGE
- TEXT
- UNITEXT
- VARBINARY

Limitations of Support

- TEXT, UNITEXT and IMAGE are supported up to 2 GB in length.
- Large objects that are replicated from other databases (such as Oracle BLOB and CLOB) can be mapped to Sybase CHAR, VARCHAR, BINARY, and VARBINARY columns. To prevent Replicat from abending if the replicated large object is bigger than the size of the target column, use the DBOPTIONS parameter with the ALLOWLOBDATATRUNCATE option in the Replicat parameter file. For more information, see Reference for Oracle GoldenGate for Windows and UNIX.
- To move data to a Sybase target from a source database that permits empty LOB columns, use the DBOPTIONS parameter with the EMPTYLOBSTRING option in the Replicat parameter file. This parameter accepts a string value and prevents Replicat from setting the target column to NULL, which is not permitted by Sybase. For more information, see Reference for Oracle GoldenGate for Windows and UNIX.
- When a source table contains multiple identical rows, it can cause LOB inconsistencies in
 the target table. This occurs when the source table lacks a primary key or other unique row
 identifier. The rows are inserted by Replicat on the target, but if the LOB data is updated in
 a subsequent source operation, it will only be replicated to the first row that was inserted
 on the target.
- Do not use NOT NULL constraints on the in-row LOB column. If you want to use NOT NULL constraints, use them on the off-row LOB column.



- If you need to fetch the in-row LOB data directly from the table you must use FETCHCOLS/ FETCHMODCOLS.
- Oracle GoldenGate for Sybase 15.7 does not support the in-row LOB column replication (however, it can still push the data into the in-row LOB column at in the Replicat database).
 This means tables included in the replication cannot have any in-row LOB columns. Oracle GoldenGate will abend if any replication table includes an in-row LOB column. If you need in-row LOB support, contact Oracle Support for further information.

Money Types

- MONEY
- SMALLMONEY

Limitations of Support

Money data types are supported with no integrity loss when moving data to a target column of the same data type without involving calculations or transformation. When calculations or transformation must be performed, Oracle GoldenGate supports a maximum value of a signed long integer (32-bits).

IDENTITY Type

The IDENTITY data type is supported for replication in one direction only, but not for a bidirectional configuration.

text, image, and unitext Data Types

With the Sybase 15.7 version, the LOB text, image, and unitext data types are now supported in BATCHSQL mode. The data length of the LOB is confined to 4K. If the records that contain LOB columns and the size exceeds more than 4K, then those records are excluded from the batches and are executed one at a time. The LOB columns in are now bound, while in previous Sybase version (15.5 or 15.0) the LOBs were not bound. You can use thee older behavior by using the DBPOTIONS LEGACYLOBREPLICATION parameter. This support is only applicable to Replicat running on Sybase version 15.7 and later.

User-Defined Data Types

User-defined data types are fully supported.

Non-Supported Sybase Data Types

This section lists the Sybase data types that Oracle GoldenGate does not support.

- The TIMESTAMP data is not supported. Timestamp columns data is captured though the data cannot be applied to the Sybase timestamp column due to a database limitation. The database populates this column automatically once that corresponding row is inserted or updated. To exclude timestamp columns from being captured by Oracle GoldenGate, use the COLSEXCEPT option of the TABLE parameter. Because the system generates the timestamps, the source and target values will be different.
- The Java rowobject data type is not supported.



Supported Operations and Objects for Sybase

This section lists the data operations and database objects that Oracle GoldenGate supports.

- The extraction and replication of insert, update, and delete operations on Sybase tables that contain rows of up to 512 KB in length.
- The maximum number of columns and the maximum column size per table that is supported by the database.
- Deferred inserts, deferred indirect inserts, deferred updates, and deferred deletes. It is
 possible that the use of deferred updates could cause primary key constraint violations for
 the affected SQL on the target. If these errors occur, use the Replicat parameter
 HANDLECOLLISIONS.
- TRUNCATE TABLE if the names of the affected tables are unique across all schemas. If the
 table names are not unique across all schemas, use the IGNORETRUNCATES parameter for
 those tables to prevent Replicat from abending.
- GETTRUNCATES and IGNORETRUNCATES by Extract and Replicat.
- Data that is encrypted with a system-encrypted password.
- Array fetching during initial loads, as controlled by the FETCHBATCHSIZE parameter.
- The BATCHSQL Replicat feature on ASE 15.7 SP110 and later on the following platforms:
 - AIX
 - Linux x64
 - Sun Solaris SPARC
 - Sun Solaris x64
 - Windows x64

In certain scenarios, the CS_NUMERIC and CS_DECIMAL data types are not supported by BatchSQL because of a bug in the Sybase specific CT Library. LOB replication is supported in BatchSql mode for Sybase database version 157 SP110 onward. This will improve the LOB replication performance. It is restricted to 16384 bytes of LOB data that means if LOB data is more than 16384 bytes, the data would not be processed through BATCHSQL mode instead the mode switched to Normal.

- Limitations on Computed Columns support are as follows:
 - Fully supports persisted computed columns. The change values are present in the transaction log and can be captured to the trail.
 - You cannot use NOT NULL constraints on in-row LOB columns. If you need to use NOT NULL constraints, do so only with off-row LOB columns.
 - Tables with non-persisted computed columns, but does not capture change data for these columns because the database does not write it to the transaction log. To replicate data for non-persisted computed columns, use the FETCHCOLS or FETCHMODCOLS option of the TABLE parameter to fetch the column data from the table. Keep in mind that there can be discrepancies caused by differences in data values between when the column was changed in the database and when Extract fetches the data for the transaction record that is being processed.
 - Replicat does not apply DML to any computed column, even if the data for that column is in the trail, because the database does not permit DML on that type of column. Data



- from a source persisted computed column, or from a fetched non-persisted column, can be applied to a target column that is not a computed column.
- In an initial load, all of the data is selected directly from the source tables, not the transaction log. Therefore, in an initial load, data values for all columns, including non-persisted computed columns, gets written to the trail or sent to the target, depending on the method that is being used. As when applying change data, however, Replicat does not apply initial load data to computed columns, because the database does not permit DML on that type of column.
- Persisted computed column that is defined as a key column, an index column, or that is part of a KEYCOLS clause in a TABLE or MAP statement are not used. If a unique key or index includes a computed column and Oracle GoldenGate must use that key, the computed column will be ignored. Additionally, if a unique key or index contains a computed column and is the only unique identifier on the table, all of the columns are used except the computed column as an identifier to find the target row. Thus, the presence of a computed column in a key or index affects data integrity if the remaining columns do not enforce uniqueness. Sybase does not support non-persisted computed columns as part of a key so neither does Oracle GoldenGate.
- To support TRUNCATE TABLE, all table names should be unique across all schemas within a database. This rule applies to Extract and Replicat.
- Limitations on Automatic Heartbeat Table support are as follows:
 - Heartbeat frequency should an integer that is divisible by 60. Oracle GoldenGate heartbeat parameter frequency is accepted in minutes, although you can use in seconds. The Sybase job scheduler uses the minutes in integer not in decimal so it is converted internally to set the frequency in minutes to the nearest possible value. For example, setting this value to 65 seconds results in the frequency being set to 1 minute; 140 seconds results in the value set to 2 minutes.
 - Data truncation occurs with a Replicat abend when it exceeds more than 1500 characters for the incoming_routing_path and outgoing_routing_path of the GG_HEARTBEAT_SEED, GG_HEARTBEAT, and GG_HEARTBEAT_HISTORY tables. The incoming_routing_path and outgoing_routing_path size of these table is set to 1500 characters in ASCII and is a 500 max bytes in multibyte characters. Ensure that the incoming and outgoing routing path strings are within the specified limit.
 - Sybase job scheduler must be configured on the ASE server prior to running Oracle GoldenGate heartbeat functionality.
 - For heartbeat table functionality to operate correctly, the login user must have the replication_role, js_admin_role, js_user_role roles.

Non-Supported Operations and Objects for Sybase

This section lists the data operations and database objects that Oracle GoldenGate does not support.

- Data that is encrypted with a user-defined password.
- Extraction or replication of DDL (data definition language) operations.
- Multi-Extract configuration. Only one Extract can reserve a context to read the Sybase transaction logs.
- Because SHOWSYNTAX is supported in the DYNSQL mode, NODYNSQL is deprecated.
- Table names that contain data with an underscore followed by some characters then a space (for example, 'zzz j') is not supported. Oracle GoldenGate cannot process records

containing this type of character string with GGSCI, DEFGEN, EXTRACT, or REPLICAT. Additionally, this type of data cannot be used with Oracle GoldenGate wildcard (*). If you do have this type of data in your table name, you must drop this kind of table name from your database, and then they restart the application to process and respect Oracle GoldenGate wildcard.



Preparing the System for Oracle GoldenGate

This chapter contains the requirements for the system and database resources that support Oracle GoldenGate.

This chapter contains the following sections:

- Database Configuration
- Database User for Oracle GoldenGate Processes
- Preparing Tables for Processing
- Preparing the Transaction Logs

Database Configuration

The Extract process makes calls directly to the Sybase Replication API on a source Sybase server. The source database on this server must be configured as follows to support data capture by Oracle GoldenGate.

- Because Extract uses the Sybase LTM to read the Sybase transaction log, it cannot run
 against a database configured with Sybase Replication Server. Only one process at a time
 can reserve a context that allows it to read the transaction log on the same database.
- Sybase Multisite availability leverages the Sybase Replication Server to replicate
 transactions to a Warm Standby and a target subscription database. Oracle GoldenGate
 for Sybase cannot capture from the primary Warm Standby but can capture from the
 Multisite availability target subscription database because SAP Sybase Rep Server is not
 in control of the Transaction Log for that database.
- Set the DSQUERY variable to the server that contains the database that Oracle GoldenGate will be using.
- The Extract process must be permitted to manage the secondary log truncation point. For more information, see Initializing the Secondary Truncation Point.
- Configure the database page size to 4k, 8k,16k, 32k, or larger. To use the
 UPGRADECHECKPOINT table_name command to updated the checkpoint the target Sybase
 database, it must be configured to have a row size of more than 2K pages to be
 programmatically created; it will fail to upgrade the checkpoint if the target database is
 configured to be a 2K page size. This is valid only for Replicat.

Database User for Oracle GoldenGate Processes

Oracle GoldenGate requires a database user account. Create this account and assign privileges according to the following guidelines.

- To preserve the security of your data, and to monitor Oracle GoldenGate processing accurately, do not permit other users, applications, or processes to log on, or operate as, the Oracle GoldenGate database user.
- Create a database user that is dedicated to Oracle GoldenGate. It can be the same user for all of the Oracle GoldenGate processes that must connect to a database:

- Extract (source database)
- Replicat (target database)
- DEFGEN utility (source or target database)
- The Extract process requires permission to access the source database. Do one of the following:
 - Grant System Administrator privileges.
 - Assign a user name with replication_role. The command to grant replication role is either:

```
sp_role 'grant', replication_role, Extract_user
Or
```

use dbname grant role replication role to Extract user

Note:

Specific DDL or DML operations may require the use of both sa_role and $replication_role$.

The Replicat process requires connect and DML privileges on the target database.

Preparing Tables for Processing

The following table attributes must be addressed in an Oracle GoldenGate environment.

Topics:

- Disabling Triggers and Cascade Constraints
- Assigning Row Identifiers
- Limiting Row Changes in Tables that do not Have a Key
- Replicating Encrypted Data

Disabling Triggers and Cascade Constraints

Disable triggers, cascade delete constraints, and cascade update constraints on target Sybase tables, or alter them to ignore changes made by the Oracle GoldenGate database user. Oracle GoldenGate replicates DML that results from a trigger or cascade constraint. If the same trigger or constraint gets activated on the target table, it becomes redundant because of the replicated version, and the database returns an error. Consider the following example, where the source tables are <code>emp_src</code> and <code>salary_src</code> and the target tables are <code>emp_targ</code> and <code>salary_targ</code>.

- A delete is issued for emp src.
- 2. It cascades a delete to salary src.
- 3. Oracle GoldenGate sends both deletes to the target.
- 4. The parent delete arrives first and is applied to emp targ.
- The parent delete cascades a delete to salary targ.



- The cascaded delete from salary_src is applied to salary_targ.
- The row cannot be located because it was already deleted in step 5.

To configure Replicat to disable target triggers at the start of its database session, take the following steps:

- 1. Assign the Replicat user the replication role.
- 2. Add the following parameter statement to the root level of the Replicat parameter file.

```
SQLEXEC "set triggers off"
```

Assigning Row Identifiers

Oracle GoldenGate requires some form of unique row identifier on source and target tables to locate the correct target rows for replicated updates and deletes.

Limiting Row Changes in Tables that do not Have a Key

If a target table has no primary key or unique key, duplicate rows can exist. It is possible for Oracle GoldenGate to update or delete too many rows in the target table, causing the source and target data to go out of synchronization without error messages to alert you. To limit the number of rows that are updated, use the DBOPTIONS parameter with the LIMITROWS option in the Replicat parameter file. LIMITROWS can increase the performance of Oracle GoldenGate on the target system because only one row is processed.

Replicating Encrypted Data

Oracle GoldenGate supports columns that are encrypted with a system-encrypted password, but not columns that are encrypted with a user-defined password. Check the tables from which you want to capture data against the following Oracle GoldenGate limitations:

- The table that contains the encrypted columns must have a primary or unique key.
- Columns that use encryption cannot be part of the primary key.

Encrypted columns are encrypted in the data files and in the log, so Extract must be configured to fetch the clear-text values from the database. To trigger this fetch, use the FETCHCOLS and FETCHMODCOLS [EXCEPT] options of the Extract TABLE parameter. FETCHCOLS forces a fetch of values that are not in the log, and FETCHMODCOLS or FETCHMODCOLS [EXCEPT] forces a fetch of values that are in the logs. Used together, these parameters ensure that the encrypted columns are always fetched from the database.

The following is an example of how to configure Extract to support the encryption. In this example, the encrypted column is cardnum.

```
TABLE ab.payments, FETCHCOLS (cardnum), FETCHMODCOLS (cardnum);
```

Preparing the Transaction Logs

To capture DML operations, Oracle GoldenGate reads the online logs. To ensure the continuity and integrity of Oracle GoldenGate processing, the logs must be configured as directed in the following sections:

Topics:

Initializing the Secondary Truncation Point

- Sizing and Retaining the Logs
- Enabling Transaction Logging

Initializing the Secondary Truncation Point

Establish a secondary log truncation point on the source system prior to running the Oracle GoldenGate Extract process. Extract uses the secondary truncation point to identify data that remains to be processed.

To initialize the secondary truncation point, log on to the database as a user with sa_role privileges and then issue the following command:

```
dbcc settrunc( 'ltm', valid )
```

By default, Extract will manage the secondary truncation point once it is established. Do not permit Extract to be stopped any longer than necessary; otherwise the log could eventually fill up and the database will halt. The only way to resolve this problem is to disable the secondary truncation point and manage it outside of Oracle GoldenGate, and then purge the transaction log. Data not yet processed by Extract will be lost, and you will have to resynchronize the source and target data.

To control how the secondary truncation point is managed, use the TRANLOGOPTIONS parameter. For more information, see Reference for Oracle GoldenGate for Windows and UNIX.

Sizing and Retaining the Logs

Retain enough log data on the source system so that Extract can start again from its checkpoints after you stop it or there is an unplanned outage. Extract must have access to the log that contains the start of the oldest uncommitted unit of work, and all logs thereafter. To determine where the Extract checkpoints are, use the INFO EXTRACT command. For more information about INFO EXTRACT, see Reference for Oracle GoldenGate for Windows and UNIX .

If data that Extract needs during processing is not retained, either in online or backup logs, one of the following corrective actions might be required:

- You might need to alter Extract to capture from a later point in time for which log data is available (and accept possible data loss on the target).
- You might need to resynchronize the source and target tables, and then start the Oracle GoldenGate environment over again.

Make certain not to use backup or archive options that cause old archive files to be overwritten by new backups on the source system. New backups should be separate files with different names from older ones. This ensures that if Extract looks for a particular log, it will still exist, and it also ensures that the data is available in case it is needed for a support case.

Enabling Transaction Logging

Use the ADD TRANDATA command to mark each source table for replication. This command uses the Sybase <code>sp_setreptable</code> and <code>sp_setrepcol</code> system procedures. ADD TRANDATA is the recommended way to mark the tables, instead of using those procedures through the database interface, but the owner or the system administrator can use them if needed. For more information, see the Sybase documentation.



To mark tables for replication with ADD TRANDATA:

- 1. On the source system, run GGSCI from the Oracle GoldenGate directory.
- 2. Log into the database from GGSCI.

DBLOGIN SOURCEDB database USERID user PASSWORD xxx

Where:

- database is the name of the database.
- user is the database owner or the system administrator. You will be prompted for the
 password. This command has encryption options for the password. For more
 information, see Reference for Oracle GoldenGate for Windows and UNIX.
- xxx is the password for the associated user.
- 3. Issue the ADD TRANDATA command for each table to be marked.

ADD TRANDATA SCHEMA.TABLE LOBSNEVER | LOBSALWAYS | LOBSALWAYSNOINDEX | LOBSIFCHANGED

Where:

 LOBSNEVER | LOBSALWAYS | LOBSALWAYSNOINDEX | LOBSIFCHANGED control whether LOB data is never propagated, only propagated if changed (the default), or always propagated. The ADD TRANDATA command will overwrite the LOB replication setting that is currently set for the table.

Note:

Some ADD TRANDATA options enable the ALWAYS_REPLICATE option of $sp_setrepcol$. If a LOB column contains a NULL value, and then another column in the table gets updated (but not the LOB), that LOB will not be captured even though ALWAYS REPLICATE is enabled.



Part IX

Using Oracle GoldenGate for Teradata

With Oracle GoldenGate for Teradata, you can deliver initial load and transactional data from other supported Oracle GoldenGate sources, such as an Oracle database.

Oracle GoldenGate for Teradata supports data filtering, mapping, and transformations unless noted otherwise in this documentation.

This part describes tasks for configuring and running Oracle GoldenGate for Teradata.

Topics

- Understanding What's Supported for Teradata
 This chapter contains information on database and table features supported by Oracle GoldenGate.
- Preparing the System for Oracle GoldenGate
- Configuring Oracle GoldenGate
- Common Maintenance Tasks



Understanding What's Supported for Teradata

This chapter contains information on database and table features supported by Oracle GoldenGate.

Topics:

- Supported Teradata Data Types
- Supported Objects and Operations for Teradata
- Non-Supported Operations for Teradata

Supported Teradata Data Types

The following table shows the Teradata data types that Oracle GoldenGate supports. Any limitations or conditions that apply follow this table.

Data type	v15.x	v16.x	
BLOB	Yes	Yes	
BYTEINT	Yes	Yes	
VARBYTE	Yes	Yes	
BIGINT	Yes	Yes	
BYTEINT	Yes	Yes	
DATE	Yes	Yes	
DECIMAL - 18 and under	Yes	Yes	
DECIMAL - 19 to 38	Yes	Yes	
DOUBLE PRECISION	Yes	Yes	
FLOAT	Yes	Yes	
INTEGER	Yes	Yes	
NUMERIC - 18 and under	Yes	Yes	
NUMERIC - 19 to 38	Yes	Yes	
REAL	Yes	Yes	
SMALLIINT	Yes	Yes	
TIME	Yes	Yes	
TIMESTAMP	Yes	Yes	
INTERVAL	Yes	Yes	
INTERVAL DAY	Yes	Yes	
INTERVAL DAY TO HOUR	Yes	Yes	
INTERVAL DAY TO MINUTE	Yes	Yes	
INTERVAL DAY TO SECOND	Yes	Yes	



Data type	v15.x	v16.x
INTERVAL HOUR	Yes	Yes
INTERVAL HOUR TO MINUTE	Yes	Yes
INTERVAL HOUR TO SECOND	Yes	Yes
INTERVAL MINUTE	Yes	Yes
INTERVAL MINUTE TO SECOND	Yes	Yes
INTERVAL MONTH	Yes	Yes
INTERVAL SECOND	Yes	Yes
INTERVAL YEAR	Yes	Yes
INTERVAL YEAR TO MONTH	Yes	Yes
CHAR	Yes	Yes
CLOB	Yes	Yes
CHAR VARYING	Yes	Yes
LONG VARCHAR	Yes	Yes
VARCHAR	Yes	Yes
GRAPHIC	Yes	Yes
LONG VARGRAPHIC	Yes	Yes
VARGRAPHIC	Yes	Yes
PERIOD (DATE)	Yes	Yes
PERIOD (TIME)	Yes	Yes
PERIOD (TIMESTAMP)	Yes	Yes
UDT	Yes	Yes

- Limitations of Support for Numeric Data Types
- Limitations of Support for Single-byte Character Data Types
- · Conditions and Limitations of Support for Multi-byte Character Data
- Limitations of Support for Binary Data Types
- Limitations of Support for Large Object Data Types
- Limitations of Support for Date Data Types
- Limitations of Support for IDENTITY Data Types

Limitations of Support for Numeric Data Types

When replicating these data types from a different type of database to Teradata, truncation can occur if the source database supports a higher precision that Teradata does.

The support of range and precision for floating-point numbers depends on the host machine. In general, the precision is accurate to 16 significant digits, but you should review the database documentation to determine the expected approximations. Oracle GoldenGate rounds or truncates values that exceed the supported precision.



Limitations of Support for Single-byte Character Data Types

Single-byte character types are fully supported within a single-byte Latin character set between other databases and Teradata. A VARCHAR or CHAR column cannot have more than 32k-1 bytes. If using UTF-16, this is 16k-2 characters.

Conditions and Limitations of Support for Multi-byte Character Data

Conditions and limitations of support for multi-byte character data are as follows:

- Install Oracle GoldenGate on a Windows or Linux replication server.
- Use the Teradata ODBC driver version 12.0.0.x or later.
- Do not use filtering, mapping, and transformation for multi-byte data types.
- A CHAR or VARCHAR column cannot contain more than 32k-1 bytes. If using UTF-16, these
 columns cannot contain more than 16k-2 characters.
- Set the ODBC driver to the UTF-16 character set in the initialization file.
- When creating Replicat groups, use the NODBCHECKPOINT option with the ADD REPLICAT command. The Replicat database checkpointing feature does not support an ODBC driver that is set to the UTF-16 character set. Checkpoints will be maintained in the checkpoint file on disk.

Limitations of Support for Binary Data Types

No limitations. These data types are supported between other source databases and Teradata targets.

Limitations of Support for Large Object Data Types

The following are limitations of support for large object data types.

- To replicate large objects from other databases to Teradata, use Teradata ODBC driver version 12.0 or higher on the target system. The target must support large objects that are delivered by ODBC.
- Enable the UseNativeLOBSupport flag in the ODBC configuration file. See the Teradata ODBC documentation.

Limitations of Support for Date Data Types

The following are limitations of support for date data types:

- DATE, TIME, and TIMESTAMP are fully supported when replicated from a different type of source database to Teradata.
- TIME with TIMESZONE, TIMESTAMP with TIMEZONE, and INTERVAL are not supported from a different type of source database to Teradata.
- Oracle GoldenGate supports timestamp data from 0001/01/03:00:00:00 to 9999/12/31:23:59:59. If a timestamp is converted from GMT to local time, these limits also apply to the resulting timestamp. Depending on the timezone, conversion may add or subtract hours, which can cause the timestamp to exceed the lower or upper supported limit.



Oracle GoldenGate does not support negative dates.

Limitations of Support for IDENTITY Data Types

IDENTITY must be configured as GENERATED BY DEFAULT AS IDENTITY on the target to enable the correct value to be inserted by Replicat.

Supported Objects and Operations for Teradata

This section lists the data operations and database objects that Oracle GoldenGate supports.

- Oracle GoldenGate supports the maximum number of columns per table that is supported by the database.
- Truncating operations are supported with the use of the GETTRUNCATES parameter with Oracle GoldenGate 12.2.x and greater.
- Limitations on Automatic Heartbeat Table support are as follows:
 - The ALTER HEARTBEATTABLE command is not supported and if used is ignored.
 - The ADD HEARTBEATTABLE command with the FREQUENCY, PURGE_FREQUENCY, or RETENTION_TIME option is not supported. When any of these options are specified with the ADD HEARTBEATTABLE command, a warning is displayed that the option is ignored.
 - Since Teradata does not have any internal event/job schedulers, automatic purging of heartbeat history data does not occur. You need to explicitly delete or truncate records periodically from the heartbeat history table.

Non-Supported Operations for Teradata

This section lists the data operations that Oracle GoldenGate does not support.

- Extract (capture)
- DDL



Preparing the System for Oracle GoldenGate

This chapter contains guidelines for preparing the database and the system to support Oracle GoldenGate. This chapter contains the following sections:

Topics:

- Preparing Tables for Processing
- ODBC Configuration for Teradata
- Database User for Oracle GoldenGate Processes for Teradata

Preparing Tables for Processing

The following table attributes must be addressed in an Oracle GoldenGate environment.

- Disabling Triggers and Cascade Constraints
- Assigning Row Identifiers

Disabling Triggers and Cascade Constraints

Disable triggers, cascade delete constraints, and cascade update constraints on target Teradata tables. Oracle GoldenGate replicates DML that results from a trigger or cascade constraint. If the same trigger or constraint gets activated on the target table, it becomes redundant because of the replicated version, and the database returns an error. Consider the following example, where the source tables are <code>emp_src</code> and <code>salary_src</code> and the target tables are <code>emp_targ</code> and <code>salary_targ</code>.

- 1. A delete is issued for emp src.
- 2. It cascades a delete to salary src.
- Oracle GoldenGate sends both deletes to the target.
- The parent delete arrives first and is applied to emp targ.
- The parent delete cascades a delete to salary targ.
- 6. The cascaded delete from salary src is applied to salary targ.
- The row cannot be located because it was already deleted in step 5.

Assigning Row Identifiers

Oracle GoldenGate requires unique row identifiers on the source and target tables to locate the correct target rows for replicated updates and deletes. Source tables can have any kind of key listed in How Oracle GoldenGate Determines the Kind of Row Identifier to Use. If there is no primary key identified on a table that has fixed-length columns, the length of one of the fixed-length columns must be below 3800 bytes.

- How Oracle GoldenGate Determines the Kind of Row Identifier to Use
- Using KEYCOLS to Specify a Custom Key

How Oracle GoldenGate Determines the Kind of Row Identifier to Use

Unless a KEYCOLS clause is used in the TABLE or MAP statement, Oracle GoldenGate selects a row identifier to use in the following order of priority:

- 1. Primary key (required for tables of a Standard Edition instance).
- 2. First unique key alphanumerically that does not contain a timestamp or non-materialized computed column.
- 3. If neither of these key types exist, Oracle GoldenGate constructs a pseudokey of all columns that the database allows to be used in a unique key, excluding those that are not supported by Oracle GoldenGate in a key or those that are excluded from the Oracle GoldenGate configuration.



If there are types of keys on a table or if there are no keys at all on a table, Oracle GoldenGate logs a message to the report file. Constructing a key from all of the columns impedes the performance of Oracle GoldenGate on the source system. On the target, this key causes Replicat to use a larger, less efficient WHERE clause.

Using KEYCOLS to Specify a Custom Key

If a table does not have an applicable row identifier, or if you prefer that identifiers are not used, you can define a substitute key, providing that the table has columns that always contain unique values. You define this substitute key by including a KEYCOLS clause within the Extract TABLE parameter and the Replicat MAP parameter. The specified key overrides any existing primary or unique key that Oracle GoldenGate finds.

ODBC Configuration for Teradata

Configure ODBC on each target system including the creation of a data source name (DSN). A DSN stores information about how to connect to the database. See the *ODBC Driver for Teradata User Guide* for complete information and setup steps:

https://docs.teradata.com/search/books? filters=prodname~%2522ODBC+Driver+for+Teradata%2522&content-lang=en-US

Database User for Oracle GoldenGate Processes for Teradata

Follow these requirements for the database user for Oracle GoldenGate processes:

- Create a database user that is dedicated to Oracle GoldenGate. It can be the same user for all of the Oracle GoldenGate processes that must connect to a database:
 - Replicat (target database)
 - The DEFGEN utility (target database)
- To preserve the security of your data, and to monitor Oracle GoldenGate processing accurately, do not permit other users, applications, or processes to log on as, or operate as, the Oracle GoldenGate database user.

• For Oracle GoldenGate to replicate to a target Teradata database, grant SELECT, INSERT, UPDATE, and DELETE on all of the target tables to the Replicat database user.



Configuring Oracle GoldenGate

Learn about the prerequisites and tasks for configuring Oracle GoldenGate Replicat for Oracle TimesTen database.

Topics:

- Creating a Checkpoint Table
- Configuring Oracle GoldenGate Replicat
- Additional Oracle GoldenGate Configuration Guidelines

Creating a Checkpoint Table

Replicat maintains its checkpoints in a checkpoint table in the Teradata target database (the database where you use <code>DBLOGIN</code>). Each checkpoint is written to the checkpoint table within the Replicat transaction. Because a checkpoint either succeeds or fails with the transaction, Replicat ensures that a transaction is only applied once, even if there is a failure of the process or the database.

Use the following GGSCI command on the target system, to create the Replicat checkpoint table.

ggsci> ADD CHECKPOINTTABLE schema.chkptabl Successfully created checkpoint table schema.chkptabl

For more information about creating a checkpoint table, see Administering Oracle GoldenGate.

Configuring Oracle GoldenGate Replicat

This section highlights the basic Replicat parameters that are required for most target database types. Additional parameters may be required, see the Oracle GoldenGate installation and configuration documentation for your target database and the *Reference for Oracle GoldenGate*.

Perform these steps on the target replication server or target database system.

- Configure the Manager process according to the instructions in Administering Oracle GoldenGate.
- 2. In the Manager parameter file, use the PURGEOLDEXTRACTS parameter to control the purging of files from the local trail.
- Create a Replicat checkpoint table. There are multiple options for this purpose, see Administering Oracle GoldenGate.
- Create a Replicat group. For documentation purposes, this group is called rep.

```
ADD REPLICAT rep, EXTTRAIL remote trail, CHECKPOINTTABLE owner.table
```

Use the EXTTRAIL argument to link the Replicat group to the remote trail that you specified for the data pump on the source server.

Use the EDIT PARAMS command to create a parameter file for the Replicat group. Include
the parameters shown in Example 29-1 plus any others that apply to your database
environment.

Example 29-1 Parameters for the Replicat Group

```
-- Identify the Replicat group:

REPLICAT rep
-- Specify database login information as needed for the database:

[TARGETDB target_dsn_name,] [USERID user id[, PASSWORD pw]]
-- Specify tables for delivery:

MAP owner.source table, TARGET owner.target table;
```

Additional Oracle GoldenGate Configuration Guidelines

The following are additional considerations to make once you have installed and configured your Oracle GoldenGate environment.

- Handling Massive Update and Delete Operations
- Preventing Multiple Connections
- Performing Initial Synchronization

Handling Massive Update and Delete Operations

Operations that update or delete a large number of rows will generate discrete updates and deletes for each row on the subscriber database. This could cause a lock manager overflow on the Teradata subscriber system, and thus terminate the Replicat process.

To avoid these errors, temporarily suspend replication for these operations and then perform them manually on the source and target systems. To suspend replication, use the following command, which suspends replication for that session only. The operations of other sessions on that table are replicated normally.

```
set session override replication on;
commit;
```

Preventing Multiple Connections

By default, the Replicat processes create a new connection for catalog queries. You can prevent this extra connection by using the DBOPTIONS parameter with the NOCATALOGCONNECT option.

Performing Initial Synchronization

Perform an initial synchronization of the source and target data before using Oracle GoldenGate to transmit transactional changes for the first time to configure an initial load, see *Administering Oracle GoldenGate*.



30

Common Maintenance Tasks

This chapter contains instructions for performing some common maintenance tasks when using the Oracle GoldenGate replication solution.

Topics:

Modifying Columns of a Table

Modifying Columns of a Table

To modify columns of a table:

- Suspend activity on the source database for all tables that are linked to Oracle GoldenGate.
- Start GGSCI.
- 3. In GGSCI, issue this command for the Replicat group:

```
INFO REPLICAT group
```

- 4. On the Checkpoint Lag line, verify whether there is any Replicat lag. If needed, continue to issue INFO REPLICAT until lag is zero, which indicates that all of the data in the trail has been processed.
- **5.** Stop the Replicat group.

```
STOP REPLICAT group
```

- 6. Perform the table modifications on the target databases.
- 7. Start the Replicat process.

```
START REPLICAT group
```

8. Allow user activity to resume on all of the source tables that are linked to Oracle GoldenGate.



Part X

Using Oracle GoldenGate for Oracle TimesTen

With Oracle GoldenGate for Oracle TimesTen, you can deliver initial load and transactional data from other supported Oracle GoldenGate sources, such as an Oracle database.

Oracle GoldenGate for Oracle TimesTen supports data filtering, mapping, and transformations unless noted otherwise in this documentation. This part describes tasks for configuring and running Oracle GoldenGate for Oracle TimesTen.

Topics:

- Understanding What's Supported for Oracle TimesTen
 This chapter contains information on database and table features supported by Oracle GoldenGate.
- Preparing the System for Oracle GoldenGate



Understanding What's Supported for Oracle TimesTen

This chapter contains information on database and table features supported by Oracle GoldenGate.

Topics:

- Supported Objects and Operations for TimesTen
- Non-Supported TimesTen Data Types and Features
- Supported TimesTen Data Types
- Limitations and Non-supported Items for Oracle TimesTen

Supported Objects and Operations for TimesTen

The following objects and operations are supported:

- Oracle GoldenGate for Oracle TimesTen supports delivery of transactional DML to user tables.
- INSERT, UPDATE, DELETE, and TRUNCATE operations are supported.

Non-Supported TimesTen Data Types and Features

The INTERVAL and ROWID data types are not supported.

Supported TimesTen Data Types

The following data types are supported for delivery, unless specifically noted in the limitations that follow:

Binary Data Types

```
(binary, varbinary)
```

Character Data Types

```
(char, nchar, nvarchar2, varchar2)
```

Date and Time Data Types

```
(date, time, timestamp, tt date, tt time, tt timestamp)
```

Numeric Data Types

```
(binary_float, binary_double, double, float, tt_bigint, tt_integer,
number, real, tt smallint, tt tinyint)
```

LOB

(blob, clob, nclob)

Limitations and Non-supported Items for Oracle TimesTen

The limitations and non-supported items for Oracle TimesTen are:

- · Capture (extraction) of DML operations is not supported.
- Capture and replication of DDL (data definition language) operations is not supported.
- The support of range and precision for floating-point numbers depends on the host machine. In general, the precision is accurate to 16 significant digits, but you should review the database documentation to determine the expected approximations. Oracle GoldenGate rounds or truncates values that exceed the supported precision.
- Oracle GoldenGate supports timestamp data from 0001/01/03:00:00:00 to 9999/12/31:23:59:59. If a timestamp is converted from GMT to local time, these limits also apply to the resulting timestamp. Depending on the time zone, conversion may add or subtract hours, which can cause the timestamp to exceed the lower or upper supported limit.
- Modifying the Primary Key Column value is not supported.
- Limitations on Automatic Heartbeat Table support are as follows:
 - Oracle GoldenGate supports only Delivery on TimesTen so no mechanism is required to populate/update the heartbeat tables using any event/job schedulers.
 - The ALTER HEARTBEATTABLE command is not supported and if used is ignored.
 - The ADD HEARTBEATTABLE command with the FREQUENCY, PURGE_FREQUENCY, or RETENTION_TIME option is not supported. When any of these options are specified with the ADD HEARTBEATTABLE command, a warning is displayed that the option is ignored.
 - Since TimesTen does not have any internal event/job schedulers, automatic purging of heartbeat history tables *cannot* occur. As such, you should explicitly drop or truncate the corresponding heartbeat objects to suit your environment.



Preparing the System for Oracle GoldenGate

This chapter contains guidelines for preparing the system to support Oracle GoldenGate: **Topics:**

- Export All Variables
- Configuring the TimesTen Client ODBC Driver
- Preparing Tables for Processing
- Configuring Oracle GoldenGate

Export All Variables

Ensure that the required system environment variables are set before proceeding.

Oracle TimesTen Client Libraries path. Example:

```
export LD_LIBRARY_PATH=/installpath/tt18.1.2.2.0/lib:$LD_LIBRARY_PATH
```

TimesTen Instance path. Example:

```
export TIMESTEN HOME=/instancepath/tt181
```

Configuring the TimesTen Client ODBC Driver

The following steps show how to configure TimesTen Client ODBC driver to connect to an Oracle TimesTen target database, which is required forOracle GoldenGate connectivity. Prior to beginning these steps, ensure that you have completed the System Requirements and Preinstallation Instructions in *Installing Oracle GoldenGate* .

Configuring ODBC on Linux

Configuring ODBC on Linux

The following steps provide minimum settings required to connect Oracle GoldenGate processes. For more detailed information on the Oracle TimesTen Client and Server configuration and information, review the following Oracle TimesTen documentation:

https://docs.oracle.com/database/timesten-18.1/TTOPR/client_server.htm#TTOPR177

1. Edit the \$TIMESTEN_HOME/conf/sys.odbc.ini file.

```
vi $TIMESTEN HOME/conf/sys.odbc.ini
```

Describe the data source in the template file. In the following example, TTCS_181 is used as
the client name for which DBLOGIN and SOURCEDB and TARGETDB are used to connect to the
database.

```
[ODBC Data Sources]
  TTCS 181=TimesTen 18.1 Client Driver
```

3. Set a logical server name for TTC_SERVER and use the server DSN value for the TTC_SERVER_DSN entry. The value of the TTC_SERVER_DSN must match the database specific server DSN that exists in the database server's sys.odbc.ini file.

```
[TTCS_181]
TTC_SERVER=ttRemoteDBServer_TT_181
TTC_SERVER_DSN=ttDatabaseDSN
```

4. Edit the \$TIMESTEN_HOME/conf/sys.ttconnect.ini file with the settings described in these steps.

```
vi $TIMESTEN HOME/conf/sys.ttconnect.ini
```

5. Create an entry of the same logical server name that was used as the value of <code>TTC_SERVER</code> within the <code>sys.odbc.ini</code> file. In this example, <code>ttRemoteDBServer_TT_181</code> is the entry name. Include an optional description value, and the required <code>Network_Address</code> and <code>TCP_PORT</code> values that point to the Oracle TimesTen database server and port.

```
[ttRemoteDBServer_TT_181]
Description=TimesTen
ServerNetwork_Address=server.company.com
TCP PORT=6625
```

6. From the Oracle GoldenGate directory on the target, verify the connection settings by running GGSCI and issuing the DBLOGIN command to log into the target database.

```
DBLOGIN SOURCEDB database, USERID db_user [, PASSWORD pw [encryption options]]
```

In this example:

- SOURCEDB database specifies the new Data Source Name.
- USERID db user, PASSWORD pw are the Replicat database user profile and password.
- encryption options is optional password encryption

Preparing Tables for Processing

This section describes the table attributes you must address in an Oracle GoldenGate environment with TimesTen.

Topics:

- Disabling Triggers and Cascade Constraints
- Assigning Row Identifiers



Disabling Triggers and Cascade Constraints

Disable triggers, cascade delete constraints, and cascade update constraints on the target tables, or alter them to ignore changes made by the Oracle GoldenGate database user. Oracle GoldenGate replicates DML that results from a trigger or cascade constraint. If the same trigger or constraint gets activated on the target table, it becomes redundant because of the replicated version, and the database returns an error. Consider the following example, where the source tables are <code>emp_src</code> and <code>salary_src</code> and the target tables are <code>emp_targ</code> and <code>salary_targ</code>.

- A delete is issued for emp src.
- 2. It cascades a delete to salary src.
- 3. Oracle GoldenGate sends both deletes to the target.
- 4. The parent delete arrives first and is applied to emp targ.
- 5. The parent delete cascades a delete to salary targ.
- 6. The cascaded delete from salary src is applied to salary targ.
- 7. The row cannot be located because it was already deleted in step 5.

Assigning Row Identifiers

Oracle GoldenGate requires some form of unique row identifier on the source and target tables to locate the correct target rows for replicated updates and deletes.

Configuring Oracle GoldenGate

Learn about the prerequisites and tasks for configuring Oracle GoldenGate Replicat for Oracle TimesTen database.

Topics:

- Configuring Oracle GoldenGate Replicat
- Additional Oracle GoldenGate Configuration Guidelines

Configuring Oracle GoldenGate Replicat

This section describes the Replicat parameters that are required for most target database types. For additional parameters that may be required, see the Oracle GoldenGate installation and configuration documentation for your target database and the *Reference for Oracle GoldenGate*.

Perform these steps on the target replication server or target database.

- Configure the Manager process according to the instructions in Administering Oracle GoldenGate.
- 2. In the Manager parameter file, use the PURGEOLDEXTRACTS parameter to control the purging of files from the local trail.
- 3. Create a Replicat checkpoint table. There are multiple options for this purpose, see *Administering Oracle GoldenGate*.
- 4. Create a Replicat group. For documentation purposes, this group is called rep.



```
ADD REPLICAT rep, EXTTRAIL remote trail, CHECKPOINTTABLE owner.table
```

Use the EXTTRAIL argument to link the Replicat group to the remote trail that you specified for the data pump on the source server.

5. Use the EDIT PARAMS command to create a parameter file for the Replicat group. Include the parameters shown in Example 29-1 plus any others that apply to your database environment.

Example 32-1 Parameters for the Replicat Group

```
-- Identify the Replicat group:

REPLICAT rep
-- Specify database login information as needed for the database:

TARGETDB target_dsn_name, [USERIDALIAS, [USERID userid, PASSWORD pw]]
-- Specify tables for delivery:

MAP owner.sourcetable, TARGET owner.targettable;
```

Additional Oracle GoldenGate Configuration Guidelines

The following are additional considerations to make once you have installed and configured your Oracle GoldenGate environment.

Performing Initial Synchronization

Perform an initial synchronization of the source and target data before using Oracle GoldenGate to transmit transactional changes for the first time to configure an initial load. See Initial Synchronization in *Administering Oracle GoldenGate*.

