

ORACLE®

Oracle Press™

Oracle GoldenGate 11g Handbook

About the Author

Robert G. Freeman works for Oracle Corporation and enjoys working with the newest technologies like Oracle GoldenGate and Oracle Exadata, and exploring the latest versions of the Oracle Database software. Robert has been working with the Oracle Database product for over two decades now and hopes to be working with it for another two decades or more. Robert is the author of several Oracle Press titles, including New Features books on Oracle Database 11g, Oracle Database 10g, and Oracle 9i. Robert is also the author of the Oracle Press book *Portable DBA: Oracle*. He is also the co-author of the Oracle Press *RMAN Backup and Recovery* series for Oracle Database 11g, Oracle Database 10g, and Oracle 9i.

About the Technical Editor

Jinyu Wang is part of the Oracle Server Technology product management team. She has managed several Oracle products over the last 12 years, including Oracle XDK, Oracle XML DB, Oracle Secure Enterprise Search (SES), Oracle Text, and Oracle GoldenGate. Jinyu's current interests are in database, data replication, XML, and information retrieval technology. Jinyu frequently speaks at conferences and publishes papers in journals. She holds a number of U.S. patents. She is the author of *Oracle Database 11g: Building Oracle XML DB Applications* (McGraw-Hill/Oracle Press, 2011) and the co-author of *Oracle Database 10g XML & SQL: Design, Build & Manage XML Applications in Java, C, C++ & PL/SQL* (McGraw-Hill/Oracle Press, 2004).

ORACLE®

Oracle Press™

Oracle GoldenGate 11g Handbook

Robert G. Freeman

**Mc
Graw
Hill**
Education

New York Chicago San Francisco Athens
London Madrid Mexico City Milan
New Delhi Singapore Sydney Toronto

Cataloging-in-Publication Data is on file with the Library of Congress

McGraw-Hill Education books are available at special quantity discounts to use as premiums and sales promotions or for use in corporate training programs. To contact a representative, please visit the Contact Us pages at www.mhprofessional.com.

Oracle GoldenGate 11g Handbook

Copyright © 2013 by McGraw-Hill Education. All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of Publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. All other trademarks are the property of their respective owners, and McGraw-Hill Education makes no claim of ownership by the mention of products that contain these marks.

Screen displays of copyrighted Oracle software programs have been reproduced herein with the permission of Oracle Corporation and/or its affiliates.

Excerpts of copyrighted Oracle user documentation have been reproduced herein with the permission of Oracle Corporation and/or its affiliates.

1234567890 DOC DOC 109876543

ISBN 978-0-07-179088-8

MHID 0-07-179088-8

Sponsoring Editor

Paul Carlstroem

Editorial Supervisor

Patty Mon

Project Manager

Sheena Uprety,
Cenveo® Publisher Services

Acquisitions Coordinator

Amanda Russell

Technical Editor

Jinyu Wang

Copy Editor

Margaret Berson

Proofreader

Paul Tyler

Indexer

Ted Laux

Production Supervisor

George Anderson

Composition

Cenveo Publisher Services

Illustration

Cenveo Publisher Services

Art Director, Cover

Jeff Weeks

Cover Designer

Pattie Lee

Information has been obtained by Publisher from sources believed to be reliable. However, because of the possibility of human or mechanical error by our sources, Publisher, or others, Publisher does not guarantee the accuracy, adequacy, or completeness of any information included in this work and is not responsible for any errors or omissions or the results obtained from the use of such information.

Oracle Corporation does not make any representations or warranties as to the accuracy, adequacy, or completeness of any information contained in this Work, and is not responsible for any errors or omissions.

This book is dedicated to my wife Carrie, and my five kids: Felicia, Sarah, Jared, Jacob, and Elizabeth. Also to my newly expected little girl whom we suspect we will call Amelia, and Bennett, our fun little goldendoodle. I'm also thankful to my father who has always inspired me, and even though I might not always be good at showing it, I'm thankful to God for the good life he's given me.



Contents at a Glance

1	Introduction to Oracle GoldenGate 11g and Replication	1
2	The Oracle GoldenGate Architecture	17
3	Installing Oracle GoldenGate	37
4	Configuring Oracle GoldenGate	71
5	Implementing Oracle GoldenGate One-Way Replication	109
6	GoldenGate Multitarget and Cascading Replication	153
7	GoldenGate Multimaster Replication	161
8	Advanced Features	175
9	Oracle GoldenGate and Oracle Real Application Clusters	219
10	GoldenGate for MySQL	231
11	Introducing Oracle GoldenGate for SQL Server	253
12	Monitoring, Troubleshooting, and Performance Tuning GoldenGate	313
13	Oracle GoldenGate Monitor and GoldenGate Director	349
14	Zero-Downtime Upgrades and Migrations Using Oracle GoldenGate	363
15	Oracle GoldenGate Veridata	369
16	Oracle GoldenGate Integration Options	377
17	GoldenGate 11.2 New Features	393
	Index	403



Contents

Acknowledgments	xvii
Introduction	xix
1 Introduction to Oracle GoldenGate 11g and Replication	1
What Is Replication and Why Use It?	2
What Kinds of Replication Are There?	2
Disk-Based Replication	3
Database Data Replication	4
Database Object Replication	7
Oracle Replication: The Past	7
Oracle Snapshots and Materialized Views	8
Oracle Advanced Replication	8
Oracle Streams	8
Oracle Heterogeneous Replication	9
Oracle GoldenGate	11
What's Next?	15
2 The Oracle GoldenGate Architecture	17
Documentation and Information Please	19
Overview of Databases That Support Oracle GoldenGate	19
The Oracle GoldenGate Architecture and Processes	21
The Oracle GoldenGate Manager Process	22
The Oracle GoldenGate Extract (or Capture) Process	23
The Oracle GoldenGate Trail Files	30
The Oracle GoldenGate Data Pump Process	33
The Oracle GoldenGate Server Collector	33
The Oracle GoldenGate Apply (Replicat) Process	34
The Oracle GoldenGate Software Command Interface (GGSCI)	34
The Oracle GoldenGate Parameter Files	34
Data Replication and Latency	34
Compatibility Between Different Versions of GoldenGate	35

3	Installing Oracle GoldenGate	37
	Oracle GoldenGate Requirements	39
	Oracle GoldenGate Oracle Database Platforms Supported	39
	Oracle GoldenGate Oracle Database Versions Supported	39
	Oracle GoldenGate Memory Requirements	40
	Oracle GoldenGate Disk Space Requirements	40
	Oracle GoldenGate Network Requirements	42
	The Demonstration Environments We Use in This Book	42
	Download and Set Up Oracle VirtualBox	42
	Creating the Oracle VirtualBox Images Used in This Book	46
	Download and Install Oracle GoldenGate	62
	What Do I Need to Know Before I Download GoldenGate?	62
	How Do I Download GoldenGate?	63
	Installing Oracle GoldenGate	64
4	Configuring Oracle GoldenGate	71
	Configuring Oracle GoldenGate After the Initial Install	73
	Configuring the Operating System Environment for GoldenGate	74
	Configuring the Oracle Database GoldenGate User	76
	Configuring for DDL Replication	80
	Configuring to Support Oracle Sequences	82
	Configuring the GoldenGate Directories and Using GGSCI	83
	Information to Collect Before You Architect Your Replication Solution	84
	Information You Will Need to Collect Before You Start	84
	Preparing Your Systems for Operation	86
	Configuration of the Oracle Database for GoldenGate Operations	87
	Configure the Database for Flashback Queries	89
	GoldenGate and Primary Keys	90
	Dealing with Deferred Constraints	91
	Verify That the Database Has Supplemental Logging Enabled	92
	Configuring the GLOBALS Parameter File	96
	Configuration and Administration of the Oracle GoldenGate Manager Process	96
	Configuration of the Oracle GoldenGate Manager Process	97
	Administration of the Oracle GoldenGate Manager Processes	101
5	Implementing Oracle GoldenGate One-Way Replication	109
	Preparing for Replication	110
	Introducing the HR Schema	110
	Our Selected Topology	112
	So, What Are We Doing?	113
	Preparing the Target Database for Replication	114
	Configuring the GoldenGate Manager Process	117

Configuring the GoldenGate Extract Process	121
Creating the Extract Parameter File	121
Registering the Extract	124
Adding the Trail File to GoldenGate and Associating It with the Extract	125
Administering the Extract Process	125
Configuring the GoldenGate Data Pump	130
Creating the GoldenGate Data Pump Extract Parameter File	130
Registering the GoldenGate Data Pump Extract	131
Administering the GoldenGate Data Pump Extract Process	132
Instantiating the Target Database Using GoldenGate	136
Configure and Register the Initial-Load Extract Process	137
Configure and Register the Initial-Load Replicat Process	138
Instantiate the Target Database with the Configured GoldenGate Processes	140
Configuring the GoldenGate Replicat Process	145
Creating the Replicat Parameter File	146
Registering the Replicat Process	148
Administering the Replicat Process	148
6 GoldenGate Multitarget and Cascading Replication	153
Configuring Multitarget Replication	154
Configure the GoldenGate Manager Process on the New Target Server	154
Add the New Target Database to an Existing or New Data Pump Extract	155
Configuring Cascading Replication	159
7 GoldenGate Multimaster Replication	161
Before You Configure Bidirectional Replication	162
The Application and Bidirectional Replication	163
The Trusted Source	164
Keys Are Important (Also Keeping SKs Unique)	164
Triggers	164
Cascaded Deletes	164
Loops	165
Conflict Detection and Resolution	165
Other Bidirectional Replication Issues	168
Example of Configuring a Bidirectional Replication	169
Configure the GoldenGate Extract Processes	169
Configure the GoldenGate Data Pump Extract Processes	170
Configure the Replicat Process for Both the ORCL.HR Schemas	171
Instantiate the GOLDENGATE2.ORCL.HR Schema	172

	Create a Conflict Handler	172
	Start Replication in Both Directions	174
8	Advanced Features	175
	Advanced Configuration	176
	Checkpoint Tables	176
	Securing Your Environment	181
	Filtering Your Data	191
	Filtering Tables	191
	Filtering Rows	194
	Mapping Column Data	197
	Table-Level Column Mapping	197
	Global Mapping	202
	Executing SQL	204
	SQLEXEC as an Option for TABLE or MAP	204
	sqlxexec as a Standalone Statement	206
	Scheduling sqlxexec	206
	Configuring for RAC	207
	Where to Install	207
	How to Capture Changes	207
	Synchronizing Nodes	208
	Threads	209
	Connecting	209
	Configuring for ASM	210
	Modify the Listener	210
	Modify the tnsnames.ora File	210
	Modify the Extract Parameter File	211
	DBLOGREADER	211
	Error Handling	212
9	Oracle GoldenGate and Oracle Real Application Clusters	219
	Oracle Maximum Availability Architecture (MAA)	220
	Disaster Recovery, High Availability, and Replication	223
	Installing GoldenGate in a RAC Environment	225
	Configuring the Application Virtual IP Address	225
	Creating an Action Script	227
	Clusterware Registration	229
	Basic Administration of the GoldenGate Application	229
10	GoldenGate for MySQL	231
	Installing GoldenGate for MySQL	232
	Preparing for the Installation of GoldenGate	
	on a MySQL Database Server	232
	Installing GoldenGate for MySQL	
	on a MySQL Database Server	233

Oracle to MySQL GoldenGate Replication	233
Setting Up the MySQL Database	234
Creating the GoldenGate Schema on MySQL	235
Performing the Initial Load from Oracle to MySQL	236
Creating the Definition File	237
Initial Loading of the MySQL Schema	238
Creating the Oracle Extract Process and the MySQL Replicat Process	240
Replicating from MySQL to Oracle	244
Defining the Objects and Generating the Definitions	244
Initial Population of the Oracle Table from MySQL Using GoldenGate	246
Configuring Replication Between MySQL and Oracle	248
11 Introducing Oracle GoldenGate for SQL Server	253
Installing and Configuring Oracle GoldenGate for SQL Server	255
Hardware Requirements	255
Windows Server Requirements	257
SQL Server Requirements and Supported Features	259
Target Database SQL Server Instance Configuration	267
Installing Oracle GoldenGate for SQL Server on Windows Configuration	270
Configuring the Manager Process	274
Manager Parameters	277
Managing Oracle GoldenGate for SQL Server	279
Database Management Considerations with Oracle GoldenGate for SQL Server	279
Configuring Basic GoldenGate Replication for SQL Server	283
Advanced GoldenGate for SQL Server	306
Configuring Oracle GoldenGate in a Windows Clustering Environment	308
Maintaining the Transaction Log Secondary Truncation Point When Extract Is Stopped	310
IDENTITY Property on Columns in Multimaster Configurations	310
Making DDL Changes to Existing Tables Configured for OGG	311
The ALTARCHIVELOGDEST Extract Parameter	312
12 Monitoring, Troubleshooting, and Performance Tuning GoldenGate	313
Monitoring Oracle GoldenGate	314
GUI-Based Tools	314
GGSCI	315
Troubleshooting	326
Common Methodology	327
Commonly Used Diagnostic Tools	327

	Common Extract Performance and Error Issues	328
	Common Replicat Issues	335
	Hung Processes	338
	Taking It Further	338
	Performance Tuning	341
	Tuning the Initial Load	341
	GoldenGate Lag	342
	Before You Begin	344
	Extract Lag	345
	Pump Lag	345
	Replicat Lag	346
	Parallelization	347
13	Oracle GoldenGate Monitor and GoldenGate Director	349
	Oracle GoldenGate Monitor	350
	Oracle GoldenGate Monitor Architecture	351
	Oracle GoldenGate Instance Prep	353
	Using Oracle GoldenGate Monitor Server	354
	Oracle GoldenGate Director	355
	The GoldenGate Director Architecture	356
14	Zero-Downtime Upgrades and Migrations Using Oracle GoldenGate	363
	Zero Downtime Methodology	364
	Failback Option	366
15	Oracle GoldenGate Veridata	369
	Veridata Functionality	370
	Configuring and Using Oracle GoldenGate Veridata	372
	Oracle GoldenGate Veridata Web	373
	Running Comparisons on the Command Line	374
16	Oracle GoldenGate Integration Options	377
	Using GoldenGate to Create a File for a Database Utility to Use	378
	Oracle GoldenGate for Flat File	379
	Installing Oracle GoldenGate Flat File	379
	The Oracle GoldenGate Flat File Infrastructure	380
	Types of Flat Files That Can Be Created	380
	DSV Files	380
	LDV Files	381
	The User Exit Properties File	381
	The Writer Process	382
	Defining the Writers	383
	Defining the Writer Mode	383
	Defining the Output File Name	384

Defining Rollover Attributes	384
Other Writer Output File Property Settings	385
Configuring File Data Content Properties	386
Including Meta File in Your Datafile Output	387
Including Statistics in Your Datafile Output	388
Including Statistics in Your Datafile Output	389
Other API Interfaces into GoldenGate	391
17 GoldenGate 11.2 New Features	393
New and Expanded Support for Database Platforms	394
Globalization	394
Security	395
Intelligent Conflict Resolution	395
Network Performance Improvements	399
Integrated Capture	399
GoldenGate and Compressed Data	399
GoldenGate and Extract Memory	400
Downstream Capture	400
Finally: A RAC-Related Change	401
Index	403



Acknowledgments

Writing a book is always an interesting endeavor. You learn new things, you experience new things in your life while you are writing it, and you try things that sometimes work and sometimes don't work. Of all the books I've written to date (I've lost count but it's something beyond 15 at this point), I've never had my life be more unsettled, had a book project be more complex and painful, and had such a hard time writing the acknowledgments and the introduction.

It was my intent when I started this project so long ago that this would be THE guide to Oracle GoldenGate. I wanted it to be more than just an introduction; I wanted it to be a book that you referenced all the time, with pages marked and dog-eared and worn from use. I want all my books to be that way. Even though this book has been a challenge, I believe you will find that the contents within do, in the end, meet my initial objective: This is THE guide to Oracle GoldenGate.

As such, there are many people to acknowledge, or this book would never have been written and would not be the book it is. A book this comprehensive and aggressive requires a lot of help. I've had help from so many people. So, special thanks go to the following. First, thanks most of all to the additional contributors to this work. This is a long list that includes (alphabetically): Mack Bell, Scott Black, Dennis Heisler, Venkatesh Kalipi, Paul Longhurst, Gene Patton, Amardeep Sidhu, and Eric Yen. All of these folks provided content for some portion of this book in varying amounts. If you read something you like, then they probably had some part in it. If you read something you don't like, I am probably the one who wrote that part.

Beyond contributors, there are those people in your life who just make life easier and make the writing of a book easier. The biggest credit of all goes to my wife Carrie. We got married during the course of the writing of this book and she deserves a lot

of credit and probably thinks she is a book widow already (especially since I'm starting my next book on the heels of this one). Bennett, my new friend (a little white-haired goldendoodle) also deserves thanks. He's a friendly two-year-old jumping furry ball of fun.

Also, thanks to my kids, to whom I dedicated this book in part. They have always been my inspiration and support. As I look back at my life, they are the best of what my life has been.

To my technical editors, Gene Patton in the beginning and then Jinyu Wang who did the technical editing for a majority of this book, I am thankful. You caught mistakes, typos, and errors, and you made this book much better. Being a technical editor is a largely thankless job. It does not pay well, and your eyes start to bug out after reading tons of material and testing examples. Gene and Jinyu deserve a lot of thanks.

To the tons of folks at Oracle Press, I owe lots of thanks, from Paul Carlstroem, to Wendy Rinaldi, Amanda Russell, and Ryan Willard.

In writing these acknowledgments, I hope I didn't leave anyone out. I don't believe I have, but I apologize if I've missed anyone and I hope you will forgive me.



Introduction

When I think about databases, sharing data, and really analyzing enterprise data, I sometimes think of the *Rime of the Ancient Mariner* by Samuel Taylor Coleridge. I had to learn parts of this poem in junior high (which was no small feat for my little brain). As I thought about what to write in this introduction, the first part of the *Rime* came to me. I include it here, for it seems to apply to the dilemma of data, data sharing, and the question of distributed data. The first part of the *Rime* goes like this:

“Day after day, day after day,
We stuck, nor breath nor motion;
As idle as a painted ship
Upon a painted ocean.

Water, water, every where,
And all the boards did shrink;
Water, water, every where,
Nor any drop to drink.”

The water to me can also represent our oceans of data, and it's everywhere. Enterprises are adrift in data, stuck in it as the hours progress. Data in different places, different platforms... data that can't truly be drunk because so often it sits inaccessible in the middle of an ocean (or if you prefer, alone on its own island) and because of its sheer volume. Day after day the data collects, and yet, we could do more with it if we were just free of the constraints that confine that data.

These constraints include physical constraints such as different database vendors with their various impedances, silos (or islands if you prefer) of data, the sheer volume

of data, and distances between various data centers. Logical constraints exist, too. These include differences in how data is logically stored (such as in tables), what the data actually means (for example, in one system the volume of some item might be defined using liters, and the definition of the volume in another system might be in gallons).

There are other problems to solve, real-time data processing and data warehousing, high availability, and disaster recovery. Data swirls around all of these problems. Indeed, as with the mariner, these problems become the albatross hung around our neck, weighing us down. Often this albatross is of our own doing, a result of not thinking and planning ahead. How do we remove this albatross from our necks? How do we stop the grinding death that confronts us as data continues to grow, and how do we figure out how to make the best use of our data? Maybe we can avoid the fate of the mariner—maybe we can find a solution to solve that problem we have created for ourselves and avoid wandering the earth telling our story.

It is our hope that this book will help you in your effort to remove the albatross from around your neck. This book is about a product called Oracle GoldenGate. Oracle GoldenGate can be a very comprehensive solution to the data problems that today's enterprises face. Oracle GoldenGate offers us the ability to collect our islands of data from the still waters and pool them into places where that data can be of even greater use and effect.

With Oracle GoldenGate, you can move data between Oracle and other databases such as DB2 or Microsoft SQL Server. You can create combined data sources (that is, operational data stores), which allow you to store, aggregate, and analyze data in different ways. By bringing all this diverse data together and discovering relationships that were previously obscured, you can develop powerful solutions.

As you will see in this book, Oracle GoldenGate is about more than just moving data between disparate data sources. Oracle GoldenGate provides high availability solutions to protect your data. These solutions also make upgrades and migrations of your database software much easier. With Oracle GoldenGate's features, you can perform near zero-downtime migrations and upgrades (any downtime is dependent on the time it takes to switch your application over to another database).

In this book, you will find a wealth of information on Oracle GoldenGate, and we hope you will find it helps you to make better use of your data, and improve your uptime. You will find a good foundational introduction to Oracle GoldenGate in the first chapter. As we progress, we will lead you through using Oracle GoldenGate step by step, from installing and configuring the product, to configuring various forms of Oracle GoldenGate replication.

Toward the middle of the book, after you have learned how to really use Oracle GoldenGate in an Oracle environment, we will cover setting up and using Oracle GoldenGate in a variety of non-Oracle environments including Oracle MySQL and Microsoft SQL Server. These chapters demonstrate one of the more powerful features of Oracle GoldenGate, its heterogeneous nature.

In the later chapters we will cover topics such as performance tuning, troubleshooting, and monitoring of Oracle GoldenGate. We will also cover various GoldenGate utilities such as the GoldenGate Monitor, Veridata. Finally, we will discuss using Oracle GoldenGate for zero-downtime operations, other GoldenGate integration options (like reading flat files and writing out SQL*Loader compatible files), and then, last but not least, we provide a chapter that highlights the newest features released in Oracle GoldenGate Version 11.2.

Intended Audience

This book is suitable for the following readers:

- DBAs who need to replicate data across Oracle databases.
- DBAs who need to replicate data across heterogeneous database environments.
- DBAs and architects who need to replicate large volumes of data, between various databases, in a widely distributed environment.
- DBAs and architects who need to replicate data over large distances in a very efficient manner.
- Database architects who need to share data across real-time OLTP databases.
- Data warehouse architects who need to design low-latency replicated operational data repositories and stores.

This book is not for the Oracle Database beginner, though it is for the Oracle GoldenGate beginner. To get the most out of this book, you will need a basic understanding of how the Oracle Database software works. You will need to understand how to create basic SQL statements and how to navigate Oracle schemas and create schema objects. With these basic administration skills, you should be able to navigate this book easily.

I hope you enjoy this book and that it helps you in your efforts to improve your enterprise database architecture.

