

Oracle Zero Downtime Migration – Logical Migration Performance Guidelines

Guidance in configuring ZDM GoldenGate settings
and choosing the correct GoldenGate Hub shape

March 2022, Version 1.0
Copyright © 2022, Oracle and/or its affiliates
Public

Purpose statement

This document provides guidance on Oracle ZDM Logical Migration performance expectations based on Oracle GoldenGate Extract and Replicat.

Disclaimer

This document in any form, software, or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced, or distributed to anyone outside Oracle without the prior written consent of Oracle. This document is not part of your license agreement, nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remain at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

Table of contents

Purpose statement	2
Disclaimer	2
Introduction	4
Architecture and overview	4
Oracle GoldenGate Hub	5
Planning logical migration performance using ZDM	6
ZDM GoldenGate configuration guidelines	6
ZDM source database guidelines	6
ZDM target database guidelines	6
Summary	10
Appendix	11
Test cases and examples	11
Test environment	11
Example resource utilization during ZDM logical migration testing	12



Introduction

Oracle Zero Downtime Migration (ZDM) is Oracle's recommended solution to migrate Oracle Database to Oracle cloud or Exadata platform.

Zero Downtime Migration follows Oracle's Maximum Availability Architecture (MAA) best practices to minimize database downtime. The physical migration option leverages Recovery Manager (RMAN) and Oracle Data Guard, and the logical migration option uses Oracle Data Pump and Oracle GoldenGate.

This technical brief focuses on the Oracle ZDM logical migration option regarding performance using Oracle GoldenGate. Depending on available source and target cloud system resources, this paper guides you to recommended ZDM deployment settings that maximize Oracle GoldenGate performance with acceptable source and target performance impact.

Architecture and overview

The ZDM logical migration option uses Data Pump to instantiate the database, a pre-fabricated Oracle GoldenGate Microservices Hub to offload most of the replication overhead, and an Oracle GoldenGate solution to minimize any downtime to potentially zero.

In most cases, the GoldenGate Hub is used as the ZDM Server. However, having separate servers for GoldenGate and ZDM is possible.

It is recommended to have the GoldenGate Hub in close proximity to the target system. This will reduce latency and improve GoldenGate replicate performance.




Figure 2: ZDM Logical Migration Architecture

For more information about Oracle Zero Downtime Migration solution, including complete documentation, refer to <https://www.oracle.com/database/technologies/rac/zdm.html>

Oracle GoldenGate Hub

ZDM logical online migration requires an Oracle GoldenGate Microservices Hub. The GoldenGate Hub is an architecture that places the GoldenGate software on a different host than the databases being operated against. The advantage of this architecture is the isolation of most GoldenGate resources from the database source and target servers. More information about GoldenGate Microservices Architecture can be found in GoldenGate documentation [Components of Oracle GoldenGate Microservices Architecture](#)

Oracle GoldenGate replication processes used by ZDM logical migration are:

- Extract - a process that captures data against the source database. The Oracle GoldenGate Extract process interacts directly with a database logmining server to receive data changes in the form of logical change records (LCR). The Extract writes the logical change records into trail files. More information about GoldenGate Extract can be found in [About Extract](#).
- Replicat - a process that delivers data to a target database. It reads the trail file, reconstructs the DML or DDL operations, and applies them to the target database. The ZDM logical migration uses a non-integrated parallel Replicat, which uses multiple apply processes to apply replicated data in parallel. More information about GoldenGate parallel Replicat can be found in [About Parallel Replicat](#)

The components of parallel Replicat are:

- Mappers - operate in parallel to read the trail, map trail records, convert the mapped records to the Replicat LCR format, and send the LCRs to the Merger for further processing. While one Mapper maps one set of transactions, the next Mapper maps the next set of transactions. The trail information is split, and the trail file is untouched because it keeps trail information in order.
- Collater and Scheduler – the master Replicat process has two threads, Collater and Scheduler. The Collater receives mapped transactions from the Mappers and puts them back into trail order for dependency calculation. The Scheduler calculates dependencies between transactions, groups the transactions into independent batches, and sends the batches to the Appliers to be applied to the target database.
- Appliers - reorder records within a batch for array execution. The Applier applies the batch to the target database and performs error handling. It also tracks applied transactions in checkpoint tables.

To set up the Oracle GoldenGate Hub, follow the instructions in [Provisioning Oracle GoldenGate Microservices on Marketplace](#)

Deploy Oracle GoldenGate Microservices with the image "Oracle GoldenGate for Oracle – Database Migrations" using the default version and verify if the image contains the latest software version. For GoldenGate Hub OCI VM shape, please check recommendation in "**GoldenGate Hub VM Shape**" below



ORACLE
GoldenGate

Oracle GoldenGate for Oracle - Database Migrations

Get App >

Oracle GoldenGate for Oracle for use with Oracle Database Migration Tools

Oracle Cloud Infrastructure | Migration, Data Integration

Contact Listing Provider

Provide the required information and create the stack. The VM is active upon creation, and you can log in using an SSH-key.

Planning logical migration performance using ZDM

This section complements [Preparing for a Logical Database Migration](#) which will lead to guidance in setting the ZDM GOLDENGATESETTINGS_EXTRACT_PERFORMANCEPROFILE and ZDM GOLDENGATESETTINGS_REPLICAT_* parameters and correct Oracle Cloud GoldenGate Hub shape.

Source database resources

- Estimated used and free CPU resources during the migration window (FREE CPU resources is what can be leveraged for the ZDM migration service)

Target database resources

- Estimated used and free CPU resources during the migration window
- Keep Autonomous Databases autoscaling setting enabled - more resources could be used automatically when needed

Prepare business requirements and maintenance window start time and duration for the migration

ZDM GoldenGate configuration guidelines

ZDM source database guidelines

ZDM source database resources are based on GoldenGate Extract performance profiles, which is configured with the ZDM parameter GOLDENGATESETTINGS_EXTRACT_PERFORMANCEPROFILE

- Depending on the workload and database version, the source database system's CPU resources can vary from a minimum LOW of less than four vCPUs (2 OCPU) to a HIGH of up to 6 vCPUs (3 OCPU) during the migration window. To avoid application impact, choose appropriately LOW-RES or HIGH settings based on free CPU resources on the source database server.
- If there are enough FREE CPU resources on the Source system, set ZDM configuration parameter GOLDENGATESETTINGS_EXTRACT_PERFORMANCEPROFILE to HIGH
- If not or there are any resources limitation (examples in the table below), then set ZDM configuration parameter GOLDENGATESETTINGS_EXTRACT_PERFORMANCEPROFILE to LOW-RES

ZDM target database guidelines

NOTE: In this document, Target ZDM configuration parameters are logically grouped to define profiles LOW and HIGH. These profiles names referred to in the document below resolve to the set of ZDM configuration parameters as explained below

ZDM Target resources should be based on Target system activities. Depending on the workload and database version, the target database system's CPU resources can vary from a minimum LOW of less than three vCPUs (2 OCPU) to a HIGH of up to all vCPUs available during the migration window. To expedite the apply performance, choose appropriately REPLICATION APPLY profile LOW or HIGH settings based on free CPU resources on the target. If there is no activity on the target database during migration, these settings lead to best performance replication (represent profile **HIGH**) and shortest migration time, set ZDM configuration parameters:

For Oracle **Autonomous Database (profile HIGH)**:

- GOLDENGATESETTINGS_REPLICAT_MAPPARALLELISM=5,
- GOLDENGATESETTINGS_REPLICAT_APPLYPARALLELISM= 2 * target pluggable database CPU_COUNT

For all **non-Autonomous Databases (profile HIGH)**:

- GOLDENGATESETTINGS_REPLICAT_MAPPARALLELISM=5
- GOLDENGATESETTINGS_REPLICAT_MINAPPLYPARALLELISM= 4 (default)
- GOLDENGATESETTINGS_REPLICAT_MAXAPPLYPARALLELISM= 2 * target pluggable database CPU_COUNT
- If there is concurrent or other activity on the target system during migration, this leads to minimal impact on the Target system. A small set of Appliers processes lead to lower data apply rate and finally, longer migration (represent profile **LOW**), set ZDM configuration parameters:

For Oracle **Autonomous Databases (profile LOW)**:

- GOLDENGATESETTINGS_REPLICAT_MAPPARALLELISM=4, (default)
- GOLDENGATESETTINGS_REPLICAT_APPLYPARALLELISM= target pluggable database CPU_COUNT / 2 (default)

For all **non-Autonomous Databases (profile LOW)**:

- GOLDENGATESETTINGS_REPLICAT_MAPPARALLELISM=4, (default)
- GOLDENGATESETTINGS_REPLICAT_MINAPPLYPARALLELISM= 4 (default)
- GOLDENGATESETTINGS_REPLICAT_MAXAPPLYPARALLELISM= target pluggable database CPU_COUNT / 2

GoldenGate Hub VM shape guidelines

GoldenGate Hub VM Shape recommendation is determined by OCI OCPUs, general formula is:

$$\#OCPU \text{ on GG Hub} = ((\#replicat_appliers / 2) + \#vCPU_extract + \#vCPU_OS) / 2$$

- **#replicat_appliers**: number of GoldenGate Replicat Appliers configured with ZDM response file using formula explained in 'ZDM Target' section above. Use value of GOLDENGATESETTINGS_REPLICAT_APPLYPARALLELISM or GOLDENGATESETTINGS_REPLICAT_MAXAPPLYPARALLELISM
- **#vCPU_extract**: based on ZDM performance profile on source using GOLDENGATESETTINGS_EXTRACT_PERFORMANCEPROFILE map to GoldenGate Hub CPU usage from less than one vCPU (1 OCPU) to a maximum of 2.5 vCPUs (2 OCPU)
- **#vCPU_OS**: used for Operating System resources management (set to 2 vCPUs)
- As an example (details in the Resources usage summary and recommendation table below): Extract at 2.5 vCPUs (2 OCPU) and Replicat at ten vCPUs (5 OCPU) with OS requirements at two vCPUs (1 OCPU) - it's sum up to 14.5 vCPUs. (8 OCPU) The following even number is 16, so OCI shape with 8 OCPUs would match the VM size - VM shape will be OCI VMStandard2.8

- If GoldenGate Hub will be deployed on Oracle OCI, supported OCI Shapes are: VMStandard2.1, VMStandard2.4, VMStandard2.8, VMStandard2.16, and VMStandard2.24. Other than OCPU resources, except block storage, are predicted by VM shape

Resources usage summary and recommendations

ZDM GOLDENGATE SETTINGS	ZDM RESOURCES: SOURCE DB	ZDM RESOURCES: GG HUB	ZDM RESOURCES: TARGET DB
HIGH (Source)	<ul style="list-style-type: none"> • 2.5-6 vCPU • 3GB of RAM (SGA steam pool) • 185MB/sec additional I/O • 70MB/s outbound network data 	<ul style="list-style-type: none"> • vCPU: 1-2.2 • Network (IN): 65-73MB/sec • I/O: 25-95MB/sec writes • Memory: 373MB 	N/A
<p>RECOMMENDATION: GoldenGate Extract with performance profile HIGH is recommended for fastest Extract performance. The logmining server, running in the source database, will potentially consume the highest CPU, I/O, and network resources on the database server, along with increased CPU, I/O, and network resources on the GoldenGate Hub, when compared to the other performance profile settings.</p> <p>For example, during a period of catchup (the workload generated a volume of redo before Extract was started) of a Swingbench OLTP style workload using Database 12c Release 2, generating an average of 14-17MB/sec of redo, GoldenGate Extract consumed the following resources:</p> <ul style="list-style-type: none"> • Database server: Logmining server consumed an average of 2.5-6 vCPU, 3GB of RAM (SGA). 185MB/sec additional I/O and 70MB/s outbound network data. • GoldenGate Hub: running on OCI Shape StandardVM_2.8 with 8.6 GB/s network interface based on Oracle Marketplace image consumed up to 1-2.2 vCPU, 25-95MB/sec I/O, 373MB RAM, and 70MB/sec inbound network data. 			
LOW_RES (Source)	<ul style="list-style-type: none"> • vCPU: 1-4 • Network (OUT): 22-26MB/sec • I/O: 73MB/sec reads • Memory: 16.8MB (streams pool) 	<ul style="list-style-type: none"> • vCPU: 0.5-2,2 • Network (IN): 22-26MB/sec • I/O: 24-40MB/sec writes • Memory: 18.6MB 	N/A
<p>RECOMMENDATION: GoldenGate Extract with the performance profile of LOW_RES uses the least amount of resources on the database server and GoldenGate Hub and offers the slowest Extract performance.</p> <p>Using the same Swingbench OLTP style workload as above, generating 14-17MB/sec of redo, GoldenGate Extract consumed the following resources:</p> <ul style="list-style-type: none"> • Database server: consumed an average of 1-4 vCPU, 17MB of RAM (SGA). 75MB/sec additional I/O and 25MB/s outbound network data. 			

	<ul style="list-style-type: none"> GoldenGate Hub: running on OCI Shape StandardVM_2.8 with 8.6 GB/s network interface based on Oracle Marketplace image consumed up to 0.5-2 vCPU, 25-50MB/sec I/O, 17MB RAM, and 25MB/sec inbound network data. <p><i>NOTE: The performance profile LOW_RES was approximately 2.5 times slower than a performance profile of HIGH.</i></p>		
profile: HIGH (Target)	N/A	<ul style="list-style-type: none"> vCPU: 8-10 Network: in/out: 2.9/14 MB/s I/O: trail read from Mem /cached/ Memory: 2143MB 	<ul style="list-style-type: none"> vCPU(s): 9.1 I/O log: 159 MB/s
<p>RECOMMENDATION: GoldenGate Replicat with ZDM performance profile HIGH uses maximum resources on the Target DB system and significantly increases resource usage on the GoldenGate Hub. Running against Target DB Replicat number of Appliers configuration should follow provided formula (#Appliers = 2x vCPU) with the option to tune down if needed.</p> <p>For example, a 2-hour Swingbench (6 users) OLTP workload generated on Source DB and processed by Extract has been used by Replicat to apply on Oracle ADB-D v.19.11; Replicat consumed the following resources:</p> <ul style="list-style-type: none"> Target Database server with ten vCPU used 9.1 CPU(s) based on AWR report and generated I/O: 159MB/s GoldenGate Hub running on OCI Shape StandardVM_2.8 used 8-10 vCPU and network outbound; 14MB/s and 2GB of memory <p>This configuration is recommended for migration with no parallel workload running on Target System. Provide high performance with the fastest apply rate—finally, shortest migration time.</p>			
profile: LOW (Target)	N/A	<ul style="list-style-type: none"> vCPU: 2,5-2.7 Network in/out 0.9/5.1 MB/s I/O: Trail read from MEM /cached/ Memory: 1072MB 	<ul style="list-style-type: none"> vCPU: 3.3 I/O log write: 65 MB/s
<p>RECOMMENDATION: GoldenGate Replicat running ZDM performance profile LOW uses minimal resources on Target DB system and GG Hub. This profile uses a formula where #Appliers equals half of Target DB vCPUs.</p> <p>For example, the same workload used for testing above consume resources as follow:</p> <ul style="list-style-type: none"> Target Database server used 3.3 vCPU(s), and I/O write 65MB/s 			

- GoldenGate Hub used during test used 2-3 vCPU thread, network 09/5.1 MB/s in/out, and 1 GB of memory.

This configuration is recommended for users already running production workload on the Target DB system where minimal impact is required. This config provides a slower apply rate compared to HIGH and finally slower site sync between Source and Target

Summary

The goal of this paper was to provide simple guidelines to set ZDM GoldenGate tuning settings and pick the correct GoldenGate Hub resource shape depending on available FREE CPU resources on the source and target.

Guidelines were described under ZDM GoldenGate Configuration Guidelines and summarized in the table above.

Appendix

Test cases and examples

To measure Oracle GoldenGate Extract and Replicat performance and their impact on the source and target database systems, tests were conducted to measure system resources required for GoldenGate processes (Extract for the source and Replicat for the target) with and without concurrent workload. The following test cases have been run:

- Baseline - a Swingbench OLTP workload is run, and the system resource profile is captured. The same workload needs to be processed by GoldenGate, which is part of the "Catchup" phase.
- Extract Catchup without a workload- GoldenGate Extract on the source system processes previous Swingbench workload. Performance and resource profile is captured. No additional concurrent workload is happening.
- Extract catchup with the active workload - GoldenGate Extract catchup and concurrent OLTP workload occur on the source system. During the first part of the test, Extract will initially run in Catchup mode (archivelog reads) and then switch to real-time mining (redo log buffer/online redo). Performance and resource profile is once again captured.
- Replicat Catchup without active workload - for each run, GoldenGate Replicat will read and apply the changes from the Swingbench OLTP workload on the target system. There is no concurrent workload while Replicat is running on Target, which is the most common migration case.
- End to end migration:

Prerequisites: Extract runs in real-time, parallel to OLTP workload generated by Swingbench OLTP on Source database. This is done once to simulate the Oracle DataPump migration period and generate trails stored on GoldenGate Hub

Testing:

- Extract runs on Source DB in parallel to OLTP Swingbench workload during each test case, when Replicat needs to catch up (prerequisites trails) and keep in Sync Source and Target databases.
- Extract runs with different performance profiles
- Replicat runs with a different number of Replicat Mappers and Replicat Appliers responsible for parallelism operation

Test environment

- Source database is version 12.2.0.1.200114, multitenant with a single PDB
- Source database running in OCI ExaCS (Frankfurt) with 16x OCPU
- All tests are run in isolation with no other databases open on the source cluster
- Database workload is generated by Swingbench (modified to include no key or think time and has only newOrder process)
- GoldenGate Hub-hosted on OCI compute shape in the same region and AD as source database: OCI VMStandard2.8 (8 OCPU, 8.2 Gbit/s)
- GoldenGate Microservices 19.1.0.0.210228, installed from OCI Marketplace
- ZDM server version 21.2 is running on the GoldenGate Hub

- The target database is Oracle Autonomous Database on Dedicated infrastructure version 19c. with 5 OCPU
- Autonomous Database autoscaling is disabled
- Replicat uses the TP service name and a COLOCATION_TAG for connection to the target ADBD database.

Example resource utilization during ZDM logical migration testing

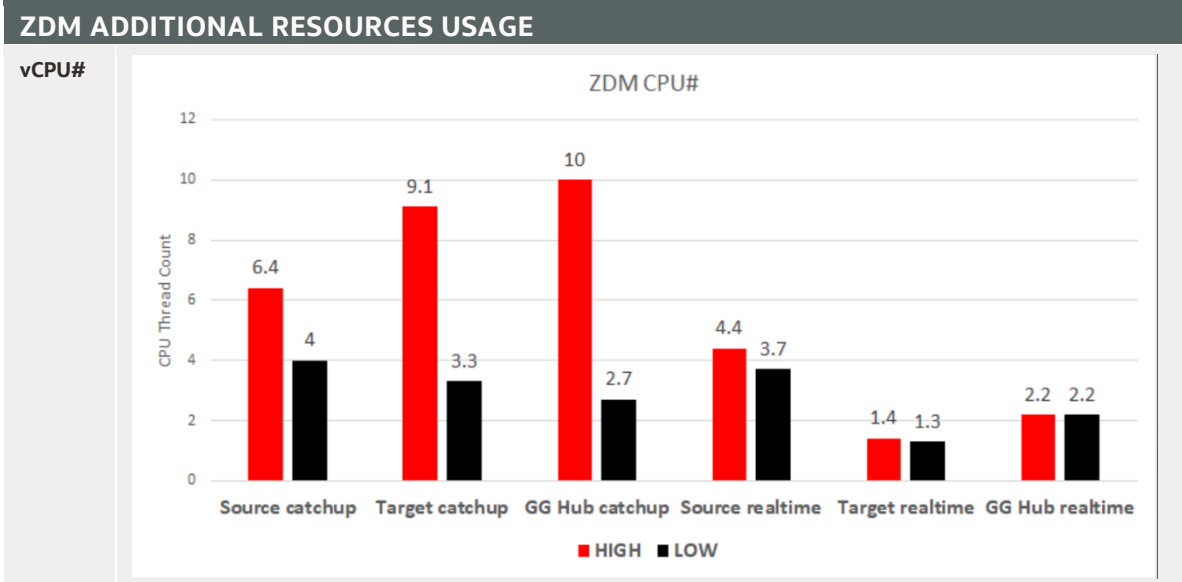
During logical migration with ZDM, there are four distinct phases:


1. Planning and prerequisites: For details, see the ZDM documentation at [Preparing for a Logical Database Migration](#)
2. Data Pump instantiation: For details, see the ZDM documentation at [Data Pump Parameter Settings for Zero Downtime Migration](#)
3. Catchup: During this phase, GoldenGate Extract and Replicat replicate changes from the source to the target. This is a period of additional workload and overhead on the source and target because of the heavy GoldenGate tasks processing transactions behind the current workload. Working as fast as possible minimizes the gap.
4. Real-time with replication: There is minimum replication lag in this period, but there is still active source workload plus GoldenGate Extract and Replicat processing for the existing workload.

Example of workload generated on the source database

WORKLOAD TYPE	SOURCE DB WORKLOAD
Two hours of OLTP Swingbench workload with six parallel users	<ul style="list-style-type: none"> • Redo logs generated: 14-17MB/s • Available CPU(s): 16 vCPU (8 OCPU) • Used CPU(s): 6 vCPU (3 OCPU) • TPS ~ 1700 • DB version: 12.2

Example of source and target databases and GoldenGate Hub additional resources usage






Example of Swingbench OLTP migration throughput based on HIGH and LOW ZDM performance profiles

The charts below show the differences in replication performance between available replication profiles HIGH and LOW for GoldenGate Extract and Replicat. To properly tune replication and limit resource usage on the

source or target system, you can mix performance profiles.



REPLICAT
LPS
records/sec



Connect with us

Call **+1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2022, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased until authorization is obtained.

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

Intel and Intel Xeon 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, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120

Disclaimer: If you are unsure whether your datasheet needs a disclaimer, read the revenue recognition policy. If you have further questions about your content and the disclaimer requirements, e-mail REVREC_US@oracle.com.