

Oracle Zero Downtime Migration – Logical Offline Migration to ExaDB-D on Oracle Database@Azure

Technical Brief

July, 2024, Version [1.0]
Copyright © 2024, Oracle and/or its affiliates
Public

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 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, timing, and pricing of any features or functionality described in this document remains 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	4
Zero Downtime Migration	5
Supported Configurations	5
Architecture	6
Zero Downtime Migration Service Host	6
Zero Downtime Migration Service Host Requirements	6
Network and Connectivity	7
Source Database	7
Target Database	8
Source and Target Database Pre-Requirements	9
Additional Configuration	9
SSH Key	9
NFS File Share via Azure Files	9
Database Migration Step by Step with ZDM	10
Step 1: Prepare the Source Database Host On-Premises	10
Step 2: Prepare the Source Database On-Premises	10
Step 3: Prepare the target database host on ExaDB-D on Oracle Database@Azure	10
Step 4: Prepare the target database on ExaDB-D on Oracle Database@Azure	11
Step 5: Prepare the ZDM Service Host On-Premises	11
Step 6: Create the Logical Offline Migration Response File on the ZDM host	12
Step 7: Evaluate the Configuration	13
Step 8: Initiate the Migration	14
Known Issues	15
Troubleshooting & Other Resources	15




Figure 1. The Oracle Zero Downtime Migration Logo comprises a Database and a Clock with an arrow pointing to a Database deployed in the Cloud.

Purpose

Oracle customers are rapidly increasing their migration of workloads into the Oracle Cloud, Engineered Systems, and Oracle Database@ Azure. However, migrating workloads has been a source of challenges for many years. Migrating database workloads from one system to another or into the Cloud is easier said than done.

Based on years of experience migrating Oracle workloads, Oracle has developed Zero Downtime Migration (ZDM). ZDM is Oracle's premier solution for a simplified and automated migration experience, providing zero to negligible downtime for the production system depending on the migration scenario. ZDM allows you to migrate your on-premises Oracle Databases directly and seamlessly to and between Oracle Database@Azure and any Oracle-owned infrastructure, including Exadata Database Machine On-Premises, Exadata Cloud at Customer (ExaDB-C@C), and Oracle Cloud Infrastructure. Oracle ZDM supports a wide range of Oracle Database versions and, as the name implies, ensures minimal to no production database impact during the migration.

ZDM follows Oracle Maximum Availability Architecture (MAA) principles¹ and incorporates products such as GoldenGate and Data Guard to ensure High Availability and an online migration workflow that leverages technologies such as the Recovery Manager, Data Pump, and Database Links.

This technical brief is a step-by-step guide for migrating your on-premises Oracle Databases to Oracle Exadata Database Service on Dedicated Infrastructure (ExaDB-D) on Oracle Database@Azure, with ZDM's Logical Offline workflow.

Oracle ZDM will run on a separate node and connect to Source and Target to perform the migration. This guide will cover all requirements for installing the Oracle ZDM Service Host, the Source Database, the Target Database recipient of the migration process, and the networking used. The migration process will be dissected and done in a step-by-step fashion. This guide will answer the most frequently asked questions regarding the product and the overall migration process.

For more information on Oracle Zero Downtime Migration, please visit ZDM's product website and Oracle Database@Azure product website.²

¹ <https://oracle.com/goto/maa>

² <https://www.oracle.com/goto/zdm>

<https://www.oracle.com/cloud/azure/oracle-database-at-azure/>

4 Oracle Zero Downtime Migration – Logical Offline Migration to ExaDB-D on Oracle Database@Azure / Version [1.0]

Zero Downtime Migration

Oracle Zero Downtime Migration (ZDM) is the Oracle Maximum Availability Architecture (MAA)-recommended solution to migrate Oracle Databases to the Oracle Cloud. ZDM's inherent design keeps in mind the migration process as straightforward as possible to ensure the most negligible impact on production workloads. The Source Database to be migrated can be on-premises, deployed on Oracle Cloud Infrastructure, or a 3rd Party Cloud. The Target Database deployment can be in Oracle Exadata Database Service on Dedicated Infrastructure (ExaDB-D) on Oracle Database@Azure, Database Cloud Service on Oracle Cloud Infrastructure (OCI) Virtual Machine, Exadata Cloud Service, Exadata Cloud at Customer, and Autonomous Database. ZDM automates the entire migration process, reducing the chance of human errors. ZDM leverages Oracle Database-integrated high availability (HA) technologies such as Oracle Data Guard and GoldenGate and follows all MAA best practices that ensure no significant downtime of production environments. Oracle ZDM supports both Physical and Logical Migration workflows. This technical brief covers a step-by-step guide for the Logical Offline Migration Workflow.

A standard Logical Offline migration will take the following steps:

1. Download and Configure ZDM.
2. ZDM Starts Database Migration.
3. ZDM Starts a Data Pump Export
4. ZDM Transfer Dump Files from the Source to the Selected Backup Location.
5. ZDM Starts a Data Pump Import Operation with Transferred Dump Files.
6. ZDM Switches Over.
7. ZDM Performs Post Migration Validations.
8. ZDM Finalizes the Migration.

Supported Configurations

Oracle ZDM supports Oracle Database versions 11.2.0.4, 12.1.0.2, 12.2.0.1, 18c, 19c and 21c. ZDM's physical migration workflow requires the Source and Target Databases to be in the same database release.

Oracle ZDM supports Source Oracle Databases hosted on Linux, Solaris, and AIX operating systems. Oracle ZDM supports single-instance databases, Oracle RAC One Node databases, or Oracle RAC databases as sources. Oracle ZDM supports Oracle Database Enterprise & Standard Edition as Source and Target Databases.

Architecture

An architectural overview of the ZDM server, the source database on-premises, the target database on Oracle Exadata Database Service on Dedicated Infrastructure (ExaDB-D) on Oracle Database@Azure, and all networks and components required are described in the diagram below:

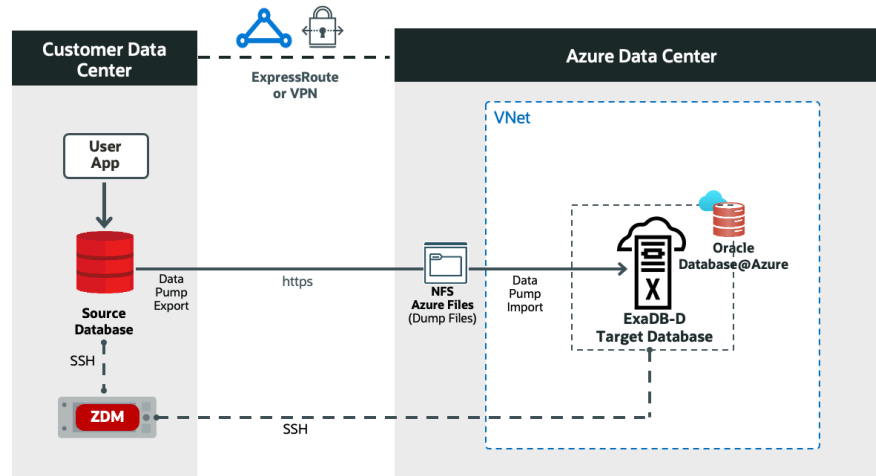


Figure 2. This is a High-Level Architectural overview showcasing the customer data center where the source database and ZDM's server reside. It also shows all connectivity to the target Oracle Exadata Database Service on Dedicated Infrastructure (ExaDB-D) on Oracle Database@Azure.

Zero Downtime Migration Service Host

Zero Downtime Migration Service Host Requirements

Oracle Zero Downtime Migration installation must take place on a separate host, which must fulfill the following requirements:

- Linux host running on Oracle 7, 8, or RHEL 8 (only these OS platforms/versions are supported).
- 100 GB of free storage space. This space is required for all the logs that ZDM will generate.
- A `zdm` group and a `zdmuser` as part of this group.
- The following packages must be installed:
 - `glibc-devel`
 - `expect`
 - `unzip`
 - `libaio`
 - `oraclelinux-developer-release-el7`
- All hostnames and IP addresses to be used must be present as entries in the `/etc/hosts` file.

For more information on the ZDM Service Host requirements and setting up ZDM on RHEL platforms, please refer to Oracle ZDM's product documentation, specifically "Setting Up Zero Downtime Migration Software" section³.

For this step-by-step guide, the ZDM Service Host runs on-premises on an Oracle Linux Server 8.9. The host private IP is masked for this guide, but as an example we will use the fictional `zz.dd.mm.hh` and the hostname is `zdmhost`.

³ <https://docs.oracle.com/en/database/oracle/zero-downtime-migration/index.html>

Network and Connectivity

Region

An Oracle Cloud Infrastructure region is a localized geographic area that contains one or more data centers, called availability domains. Regions are independent of other regions, and vast distances can separate them (across countries or continents).

Virtual Cloud Network (VCN) and subnet

A VCN is a customizable, software-defined network that you set up in an Oracle Cloud Infrastructure region. Like traditional data center networks, VCNs give you complete control over your network environment. A VCN can have multiple non-overlapping CIDR blocks that you can change after you create the VCN. You can segment a VCN into subnets, which can be scoped to a region or an availability domain. Each subnet consists of a contiguous range of addresses that don't overlap with the other subnets in the VCN. You can change the size of a subnet after creation. A subnet can be public or private.

OCI Network Security Group (NSG)

A network security group (NSG) provides a virtual firewall for cloud resources with the same security posture. For example, a group of compute instances performs the same tasks and thus needs to use the same set of ports.

Azure VNet

Azure Virtual Network (VNet) is the fundamental building block for your private network in Azure. VNet enables many Azure resources, such as Azure virtual machines (VM), to securely communicate with each other, the internet, and on-premises networks.

Azure Delegated Subnet

Subnet delegation is Microsoft's ability to inject a managed service, specifically a platform-as-a-service service, directly into your virtual network. This means you can designate or delegate a subnet to be a home for an externally managed service inside your virtual network. In other words, that external service will act as a virtual network resource, even though technically it is an external platform-as-a-service service.

Virtual network interface card (vNIC)

The services in Azure data centers have physical network interface cards (NICs). Virtual machine instances communicate using virtual NICs (vNICs) associated with the physical NICs. Each instance has a primary vNIC that's automatically created and attached during launch and is available during the instance's lifetime.

Azure Route table (User Defined Route – UDR)

Virtual route tables contain rules to route traffic from subnets to destinations outside a VNet, typically through gateways. Route tables are associated with subnets in a VNet.

Source Database

The source database runs on-premises on an Oracle Linux Server 7.7 for this step-by-step guide. The host's private IP is masked for this guide, but as an example, we will use the fictional aa.bb.sr.db address and the hostname is **onphost**. The source Oracle database is a single-instance Enterprise Edition database version 19.21 with multitenant architecture. The database name is **oradb**, and its unique name is **oradb_onp**.

The HR schema to be migrated resides in the source PDB **pdbsrc**.

Target Database

Oracle Database@Azure offers the following products:

- **Oracle Exadata Database Service on Dedicated Infrastructure**
 - You can provision flexible Exadata systems that allow you to add database compute servers and storage servers to your system anytime after provisioning.
- **Oracle Autonomous Database Serverless**
 - Autonomous Database provides an easy-to-use, fully autonomous database that scales elastically, delivers fast query performance, and requires no database administration.

Oracle Database@Azure integrates Oracle Exadata Database Service, Oracle Real Application Clusters (Oracle RAC), and Oracle Data Guard technologies into the Azure platform. The Oracle Database service runs on Oracle Cloud Infrastructure (OCI) and is co-located in Microsoft Azure data centers. The service offers features and price parity with OCI. Users purchase the service on Azure Marketplace.

Oracle Database@Azure service offers the same low latency as other Azure-native services and meets mission-critical workloads and cloud-native development needs. Users manage the service on the Azure console and with Azure automation tools. The service is deployed in Azure Virtual Network (VNet) and integrated with the Azure identity and access management system. The OCI and Oracle Database metrics and audit logs are natively available in Azure. The service requires that users have an Azure tenancy and an OCI tenancy.




Figure 3 Network and Architectural Overview

For this step-by-step guide, the target platform is Oracle Exadata Database on Dedicated Infrastructure (ExaDB-D) on Oracle Database@Azure. The infrastructure contains a 2-node VM cluster. The VM cluster hosts private IPs are not provided for this guide, but as an example, we will use the fictional ta.db.oa.1 and ta.db.oa.2, and the host names are exadbazure1 and exadbazure2. ZDM requires configuring a database target environment before beginning the migration process. The target Oracle database is a 2-node Oracle RAC version 19.22 with multitenant architecture created using Oracle Cloud Console. The database name is **oradb**, and the database's unique name is **oradb_exa**.

The HR schema is to be migrated to the target PDB **pdbtgt**.

Source and Target Database Pre-Requisites

- The character set on the source database must be the same as the target database.
- The DATAPUMP_EXP_FULL_DATABASE and DATAPUMP_IMP_FULL_DATABASE roles are required.
- The DATAPUMP_IMP_FULL_DATABASE role is required for the import operation at the specified target database for the specified target database user.

Additional Configuration

SSH Key

Check the key format:

```
[zdmuser@zdmhost ~]$ head -n1 id_rsa
```

Create an SSH Key in RSA format (if not already created):

```
[zdmuser@zdmhost ~]$ ssh-keygen -m PEM -t rsa
```

Change an existing SSH key into RSA format (if already created and need to reformat):

```
[zdmuser@zdmhost ~]$ ssh-keygen -p -m PEM -f id_rsa
```

NFS File Share via Azure Files

ZDM Logical Offline migration workflow uses Oracle Data Pump export and import to migrate the data from the source to the target database. An NFS file share is provided through the Azure Files service to store the Data Pump dump files. For this step-by-step guide, the file share path is **/azurefilesnfs/dumpfiles/**. The NFS share must be mounted on the source and target database hosts.

The NFS-mounted path is not readable for the target database user unless the Unique Identifiers (UIDs) for the source database user match with the target database user (see Oracle Zero Downtime Migration Product Documentation).

Example of mounting the NFS share:

```
sudo mount -t nfs odaamigration.file.core.windows.net:/odaamigration/testmigration
/azurefilesnfs -o vers=4,minorversion=1,sec=sys
```

Database Migration Step by Step with ZDM

Step 1: Prepare the Source Database Host On-Premises

Copy the SSH public key of the zdmuser from the ZDM host to the `.ssh/authorized_keys` file on the source database host for the user you want to use for login, in this case, **onpuser**:

```
#on ZDM host as zdmuser
[zdmuser@zdmhost ~]$ cat .ssh/id_rsa.pub
#on the source database host as user onpuser
[onpuser@onphost ~]$ vi .ssh/authorized_keys
#insert the public key and save the changes
```

Add the target database hostname, IP address, and SCAN name to the `/etc/hosts` file. As root user:

```
[root@onphost ~]# vi /etc/hosts
#add the following entries
ta.db.oa.1 oradb_exa_sample.oravcn.sample.com target
ta.db.oa.1 demo-scan-sample.oravcn.sample.com target-scan
```

Step 2: Prepare the Source Database On-Premises

Prepare the source database. As SYS user:

```
-- Set streams_pool_size to 2G
SQL> alter system set streams_pool_size=2G scope=both;
System altered.
SQL> grant DATAPUMP_EXP_FULL_DATABASE to system container=all;
Grant succeeded.
```

Step 3: Prepare the target database host on ExaDB-D on Oracle Database@Azure

Copy the SSH public key of the zdmuser from the ZDM host to the `.ssh/authorized_keys` file on the target database host for the user you want to use for login; in this case, **opc**:

```
#on ZDM host as zdmuser
[zdmuser@zdmhost ~]$ cat .ssh/id_rsa.pub

#on the target database hosts as user opc (on all VMs of the VM cluster)
[opc@exadbazure1 ~]$ vi .ssh/authorized_keys
#insert the public key and save the changes
[opc@exadbazure2 ~]$ vi .ssh/authorized_keys
#insert the public key and save the changes
```

#Add the source database hostname and IP information into the `/etc/hosts` file. As root user (on all VMs of the VM cluster)

```
[root@exadbazure1 ~]# vi /etc/hosts
#add the following entries
aa.bb.sr.db onphost
```

```
[root@exadbazure2 ~]# vi /etc/hosts
#add the following entries
aa.bb.sr.db onphost
```

Step 4: Prepare the target database on ExaDB-D on Oracle Database@Azure

Prepare the Target Database, as SYS user:

```
-- Set streams_pool_size to 2G
SQL> alter system set streams_pool_size=2G scope=both;
System altered.
SQL> grant DATAPUMP_IMP_FULL_DATABASE to system container=all;
Grant succeeded.
```

Step 5: Prepare the ZDM Service Host On-Premises

Add the source and target hostnames and IP addresses into the `/etc/hosts` file. As root user:

```
[root@zdmhost ~]# vi /etc/hosts
#add the following entries
ta.db.oa.1 exadbazure1
ta.db.oa.2 exadbazure2
aa.bb.sr.db onphost
```

Test the SSH connectivity to the source and target database hosts:

```
[zdmuser@zdmhost ~]$ ssh -i /home/zdmuser/.ssh/id_rsa onpuser@onphost
[zdmuser@zdmhost ~]$ ssh -i /home/zdmuser/.ssh/id_rsa opc@exadbazure1
[zdmuser@zdmhost ~]$ ssh -i /home/zdmuser/.ssh/id_rsa opc@exadbazure2
```

Verify that TTY is disabled for the SSH-privileged user. If TTY is disabled, the following command returns the date from the remote host without any errors:

```
[zdmuser@zdmhost ~]$ ssh -oStrictHostKeyChecking=no -i /home/zdmuser/.ssh/id_rsa
onpuser@onphost "/usr/bin/sudo /bin/sh -c date"
[zdmuser@zdmhost ~]$ ssh -oStrictHostKeyChecking=no -i /home/zdmuser/.ssh/id_rsa
opc@exadbazure1 "/usr/bin/sudo /bin/sh -c date"
[zdmuser@zdmhost ~]$ ssh -oStrictHostKeyChecking=no -i /home/zdmuser/.ssh/id_rsa
opc@exadbazure2 "/usr/bin/sudo /bin/sh -c date"
```

These commands should execute without any prompting and return the date from the remote host.

Step 6: Create the Logical Offline Migration Response File on the ZDM host

You'll find a template on the ZDM host at `$ZDMHOME/rhp/zdm/template/zdm_template.rsp`, briefly describing the parameters and their possible values. Here, we will create a new response file with the minimal parameters required.

As `zdmuser`:

```
[zdmuser@zdmhost ~]$ vi /home/zdmuser/logical_offline/logical_offline.rsp
#add the following parameters and save the changes
# migration method
MIGRATION_METHOD=OFFLINE_LOGICAL
DATA_TRANSFER_MEDIUM=NFS

# data pump
DATAPUMPSETTINGS_JOBMODE=SCHEMA
INCLUDEOBJECTS-1=owner:HR
DATAPUMPSETTINGS_METADATAREMAPS-
1=type:REMAP_TABLESPACE,oldValue:USERS,newValue:DATA
DATAPUMPSETTINGS_DATAPUMPPARAMETERS_EXPORTPARALLELISMDEGREE=2
DATAPUMPSETTINGS_DATAPUMPPARAMETERS_IMPORTPARALLELISMDEGREE=2
DATAPUMPSETTINGS_EXPORTDIRECTORYOBJECT_NAME=DUMP_DIR
DATAPUMPSETTINGS_IMPORTDIRECTORYOBJECT_NAME=DUMP_DIR

# on source and target db: select directory_name, directory_path from
dba_directories;
DATAPUMPSETTINGS_EXPORTDIRECTORYOBJECT_PATH=/azurefilesnfs/dumpfiles
DATAPUMPSETTINGS_IMPORTDIRECTORYOBJECT_PATH=/azurefilesnfs/dumpfiles

# source db (pdb)
SOURCEDATABASE_CONNECTIONDETAILS_HOST=onphost
SOURCEDATABASE_CONNECTIONDETAILS_PORT=1521
SOURCEDATABASE_CONNECTIONDETAILS_SERVICENAME=pdbsrc
SOURCEDATABASE_ADMINUSERNAME=SYSTEM

# target db (pdb)
TARGETDATABASE_CONNECTIONDETAILS_HOST=exadbazure1
TARGETDATABASE_CONNECTIONDETAILS_PORT=1521
TARGETDATABASE_CONNECTIONDETAILS_SERVICENAME=test.ocitestvm.ocitestvnt.test.com
TARGETDATABASE_ADMINUSERNAME=SYSTEM

# oci cli
OCIAUTHENTICATIONDETAILS_USERPRINCIPAL_USERID=ocidl.user.oc1..aaaaaaa
OCIAUTHENTICATIONDETAILS_USERPRINCIPAL_TENANTID=ocidl.tenancy.oc1..aaaaaaa
OCIAUTHENTICATIONDETAILS_USERPRINCIPAL_FINGERPRINT=aaa.bbb.ccc.ddd
OCIAUTHENTICATIONDETAILS_USERPRINCIPAL_PRIVATEKEYFILE=/home/zdmuser/.oci/oci_api_key.pem
OCIAUTHENTICATIONDETAILS_REGIONID=us-ashburn-1
```

Though NFS is used instead of OCI Object Storage, the OCI CLI parameters are still required to discover the target database.

Step 7: Evaluate the Configuration

Execute the following command on the ZDM host as `zdmuser` to evaluate the migration. ZDM will check the source and target database configurations. The actual migration will not be started. On the ZDM host as `zdmuser`:

```
[zdmuser@zdmhost ~]$ $ZDMHOME/bin/zdmcli migrate database \
-rsp /home/zdmuser/logical_offline/logical_offline.rsp \
-sourcenode onphost \
-sourcesid oradb \
-srcauth zdmauth \
-srcarg1 user:azureuser \
-srcarg2 identity_file:/home/zdmuser/.ssh/id_rsa \
-srcarg3 sudo_location:/usr/bin/sudo \
-targetnode exadbazure1 \
-tgtauth zdmauth \
-tgtarg1 user:opc \
-tgtarg2 identity_file:/home/zdmuser/.ssh/id_rsa \
-tgtarg3 sudo_location:/usr/bin/sudo \
-eval
```

```
Enter source database administrative user "SYSTEM" password:
Enter target database administrative user "SYSTEM" password:
Operation "zdmcli migrate database" scheduled with the job ID "1".
```

If the source database uses ASM for storage management, use `-sourcedb <db_unique_name>` instead of `-sourcesid <SID>` in the `zdmcli` command. Check the job status. On the ZDM host as `zdmuser`:

```
[zdmuser@zdmhost ~]$ $ZDMHOME/bin/zdmcli query job -jobid 1
...
Job ID: 11
User: zdmuser
Client: zdmhost
Job Type: "EVAL"
...
Current status: SUCCEEDED
Result file path: "/home/zdmuser/zdm/zdmbase/chkbase/scheduled/job-1-2024-07-03-10:40:04.log"
Metrics file path: "/home/zdmuser/zdm/zdmbase/chkbase/scheduled/job-1-2024-07-03-10:40:04.json"
...
ZDM_VALIDATE_TGT ..... COMPLETED
ZDM_VALIDATE_SRC ..... COMPLETED
ZDM_SETUP_SRC ..... COMPLETED
ZDM_PRE_MIGRATION_ADVISOR ..... COMPLETED
ZDM_VALIDATE_DATAPUMP_SETTINGS_SRC .... COMPLETED
ZDM_VALIDATE_DATAPUMP_SETTINGS_TGT .... COMPLETED
ZDM_PREPARE_DATAPUMP_SRC ..... COMPLETED
ZDM_DATAPUMP_ESTIMATE_SRC ..... COMPLETED
ZDM_CLEANUP_SRC ..... COMPLETED
```

Detailed information about the migration process can be found by monitoring the log file:

```
[zdmuser@zdmhost ~]$ tail -f /home/zdmuser/zdm/zdmbase/chkbase/scheduled/job-1-2024-07-03-10:40:04.log
```

In case troubleshooting is required, please check the ZDM server log on the ZDM Service Host under the following location:

```
$ZDM_BASE/crsdata/<zdm_service_host>/rhp/zdmserver.log.0
```

Step 8: Initiate the Migration

To initiate the actual migration, execute the same command for evaluation, but this time without the `-eval` parameter.

On the ZDM host as `zdmuser`:

```
[zdmuser@zdmhost ~]$ $ZDMHOME/bin/zdmcli migrate database \
-rsp /home/zdmuser/logical_offline/logical_offline.rsp \
-sourcenode onphost \
-sourcesid oradb \
-srcauth zdmauth \
-srcarg1 user:azureuser \
-srcarg2 identity_file:/home/zdmuser/.ssh/id_rsa \
-srcarg3 sudo_location:/usr/bin/sudo \
-targetnode exadbazure1 \
-tgtauth zdmauth \
-tgtarg1 user:opc \
-tgtarg2 identity_file:/home/zdmuser/.ssh/id_rsa \
-tgtarg3 sudo_location:/usr/bin/sudo
Enter source database administrative user "SYSTEM" password:
Enter target database administrative user "SYSTEM" password:
Operation "zdmcli migrate database" scheduled with the job ID "2".
```

Check the job status. On the ZDM host as `zdmuser`:

```
[zdmuser@zdmhost ~]$ $ZDMHOME/bin/zdmcli query job -jobid 2
...
Job ID: 2
User: zdmuser
Client: zdmhost
Job Type: "MIGRATE"
...
Current status: PAUSED
Result file path: "/home/zdmuser/zdm/zdmbase/chkbase/scheduled/job-2-2024-07-3-10:52:03.log"
Metrics file path: "/home/zdmuser/zdm/zdmbase/chkbase/scheduled/job-2-2024-07-3-10:52:03.json"
...
ZDM_VALIDATE_TGT ..... COMPLETED
ZDM_VALIDATE_SRC ..... COMPLETED
ZDM_SETUP_SRC ..... COMPLETED
ZDM_PRE_MIGRATION_ADVISOR ..... COMPLETED
ZDM_VALIDATE_DATAPUMP_SETTINGS_SRC .... COMPLETED
ZDM_VALIDATE_DATAPUMP_SETTINGS_TGT .... COMPLETED
ZDM_PREPARE_DATAPUMP_SRC ..... COMPLETED
ZDM_DATAPUMP_ESTIMATE_SRC ..... COMPLETED
ZDM_PREPARE_DATAPUMP_TGT ..... COMPLETED
ZDM_DATAPUMP_EXPORT_SRC ..... COMPLETED
ZDM_TRANSFER_DUMPS_SRC ..... COMPLETED
ZDM_DATAPUMP_IMPORT_TGT ..... COMPLETED
ZDM_POST_DATAPUMP_SRC ..... COMPLETED
ZDM_POST_DATAPUMP_TGT ..... COMPLETED
ZDM_POST_ACTIONS ..... COMPLETED
ZDM_CLEANUP_SRC ..... COMPLETED
```

Detailed information about the migration process can be found by monitoring the log file:

```
[zdmuser@zdmhost ~]$ tail -f /home/zdmuser/zdm/zdmbase/chkbase/scheduled/job-2-2024-07-3-10:52:03.log
```

Known Issues

All common issues are documented and updated periodically in Oracle Zero Downtime Migration's documentation, specifically on the product release note, Known Issues section:

<https://docs.oracle.com/en/database/oracle/zero-downtime-migration/>

Troubleshooting & Other Resources

For Oracle ZDM log review:

- ZDM Server Logs:
 - o Check - `$ZDM_BASE/crsdata/<zdm_service_node>/rhp/rhpserver.log.0`
- Check source node logs
 - o - `<oracle_base>/zdm/zdm_<src_db_name>_<job_id>/zdm/log`
- Check target node logs.
 - o - `<oracle_base>/zdm/zdm_<tgt_db_name>_<job_id>/zdm/log`

For all Oracle Support Service Requests related to Zero Downtime Migration, please be sure to follow the instructions in My Oracle Support Document:

- SRDC – Data Collection for Database Migration Using Zero Downtime Migration (ZDM) (DOC ID 2595205.1)
- <https://support.oracle.com/epmos/faces/DocContentDisplay?id=2595205.1>

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 © 2024, Oracle and/or its affiliates. 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.

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