



ORACLE

# Oracle Data Guard 19c and SAP NetWeaver® Application Server ABAP/Java

A Complementary Guide for Oracle Exadata Database Service

September 2025, Version 1.0

Copyright © 2025, 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.

## Revision History

The following revisions have been made to this paper.

Date	Revision
September 2025	Initial publication

## Table of Contents

---

<b>Scope and Assumptions</b>	<b>5</b>
<b>Terminology</b>	<b>5</b>
<b>Planning</b>	<b>7</b>
Scenarios and Complexity	7
Planning the Virtual Cloud Networks	9
Planning the Deployment of the VM Clusters	10
<b>Preparing the Primary Database</b>	<b>11</b>
Enable Force Logging on Primary Database	13
Set File Name Conversion Parameters	13
Set Log Archive Parameters	13
Set Oracle Data Guard Broker–Specific Parameters	13
Set the SERVICE_NAMES Parameter	14
Add Standby Redo Logs to the Primary Database	14
Move the Database Password File to ASM	14
Move the TDE Wallet	15
Create New Archive Directory	15
Configure Database as Primary in Oracle Clusterware	15
Restart Database Instances on the Primary Using srvctl	15
Set Up TNSNAMES.ORA for RMAN DUPLICATE	15
Final Checks on the Primary	16
<b>Setting Up the Standby Site</b>	<b>16</b>
Stop the SAP Database on Standby and Remove It from Oracle Clusterware	17
Remove the SAP Database from ASM	17
Create TDE Wallet and Archive Directories on ASM	18
Copy the TDE Wallet, the Autologin Wallet, and the Password File from Primary	18
Configure TNSNAMES.ORA	19
Configure and Start the Listener for RMAN DUPLICATE	19
Create an Oracle Initialization Parameter File for RMAN DUPLICATE	20
Start the Database Instance for RMAN DUPLICATE	20
Run RMAN DUPLICATE	21
Set Initialization Parameters for Standby	22
Configure the Database in Clusterware and Start It	22
Configure the Primary and Standby with Data Guard	22
Add SAP-Specific Database Services to Oracle Clusterware	23
<b>Implementing the Fast Start Failover (Optional)</b>	<b>25</b>
<b>Testing the Configuration</b>	<b>27</b>

<b>Installing SAP NetWeaver Application Server ABAP/Java</b>	<b>27</b>
<b>References</b>	<b>28</b>
SAP Notes	28
Oracle Technical Briefs	28

## Scope and Assumptions

This document serves as a *complementary* guide for customers who are running the SAP NetWeaver Application Server ABAP/Java platform on Oracle Exadata Database Service and are seeking to properly configure Oracle Data Guard within their SAP landscape.

Customers must have strictly followed the relevant technical briefs for their specific Oracle Exadata Cloud Infrastructure environment. The referenced technical briefs for installing and configuring SAP NetWeaver Application Server ABAP/Java on Oracle Exadata Database Service can be found on the [Oracle Help Center](#).

Customers migrating to Oracle Exadata Database Service should also refer to [Migrating SAP NetWeaver Based Systems Within the Scope of Oracle Databases](#).

Customers who are already using Oracle Multitenant may need to consider multitenant-specific aspects when transferring the TDE wallet and the autologin wallet. These aspects are currently not covered in this document.

For Oracle Data Guard 19c deployments on Oracle Cloud Infrastructure (OCI), this document explains how to set up a basic primary-standby database scenario across two different regions using Oracle Data Guard Broker. Additionally, and as an optional step, it describes how to configure Fast Start Failover.

Combining Oracle Exadata Cloud@Customer with Oracle Exadata Database Service is possible, but it requires different steps when configuring network connectivity between the customer’s data center and the corresponding Oracle Cloud data center. Although this scenario is not covered in this document, most of the configuration tasks are the same, and customers primarily need to address the differences in connecting the primary and standby sites.

## Terminology

The following tables define some of the terms used within the context of Oracle Exadata Database Service and this document. They also define terms related to Oracle-specific and SAP-specific work areas that you will likely interact with during deployment and while running SAP NetWeaver Application Server ABAP/Java on a VM cluster.

**Table 1: Terms Used Within the Context of Oracle Exadata Database Service and This Document**

TERM	DEFINITION
Exadata Cloud Infrastructure	The specific physical infrastructure that powers an Oracle Exadata Database Service. The infrastructure typically consists of two, four, or eight DB servers or Exadata compute nodes, a number of storage servers, plus more components for networking and power supply.
DB server	A database (DB) server is one physical node of the Exadata Cloud Infrastructure, often referred as an Exadata compute node. It is configured as a KVM-based virtualization host (Dom0) and can run multiple virtual machines (VMs).
Storage server	A storage server provides shared disk storage and runs the Exadata Storage Server Software. This software provides unique and powerful technology such as Smart Scan, Smart Flash Cache, Smart Flash Logging, IO Resource Manager, Storage Indexes, and Hybrid Columnar Compression.
VM cluster	A cluster that consists of two or more VMs running on different DB servers.
VM cluster node	A VM with all the components to run Oracle Real Application Clusters (RAC) Databases. This includes the presence of the operating system (OS), Oracle Grid Infrastructure software (Clusterware, ASM, ACFS, and so on), and Oracle Database software. A VM cluster node is also referred as an Oracle Database compute node, and might also be referred to as a VM, KVM guest, DomU, or virtual compute node.

TERM	DEFINITION
Virtual cloud network (VCN)	A virtual, private network that you set up in Oracle data centers. It closely resembles a traditional network, with firewall rules and specific types of communication gateways that you can choose to use. A VCN resides in a single OCI region and covers one or more CIDR blocks. For more information, see <a href="#">Networking Overview</a> .
VM Cluster Node Subsetting	Allows a VM cluster to span two or more DB servers (up to the maximum number of DB servers available in your Oracle Exadata Infrastructure). VM Cluster Node Subsetting lets you add new VM cluster nodes to an existing VM cluster if your workload or high availability requirements demand additional resources. VM Cluster Node Subsetting also lets you reduce the number of VM cluster nodes in a VM cluster down to a minimum of two.

**Table 2: Oracle-Specific Work Areas**

ORACLE WORK AREA	PURPOSE	RELATED NOTES AND COMMENTS
Oracle Cloud Console	Used to manage all the resources in OCI, for example, the deployment of VCNs, VM clusters for Oracle Exadata Database Service, or independent compute nodes for SAP NetWeaver Application Servers.	Register at <a href="https://cloud.oracle.com">cloud.oracle.com</a> .
Oracle Linux	Used on all VM cluster nodes. Other compute instances might run different OSs.	Oracle Linux 8 on VM cluster nodes is managed from the Oracle Cloud Console. Oracle Linux or other OSs on other compute instances must be managed by the customer. For Oracle Linux 8, see SAP Note <a href="#">2936683</a> . For Oracle Linux 9, see SAP Note <a href="#">3399081</a> .
Oracle Grid Infrastructure	Manage the HA of all database services, ACFS, and NFS by using tools like SRVCTL and CRSCTL, as well as application virtual IP addresses.	For HA-NFS, see <a href="#">My Oracle Support Note 1934030.1: ODA (Oracle Database Appliance): HowTo export ACFS (cloudfs) using HANFS</a> .
Oracle Database software	Use MOPatch or OPatch and SAP Bundle Patches to patch existing Oracle Database homes.	SAP Bundle Patch installation is documented in the readme file.
Oracle Automatic Storage Management (ASM)	Manage ASM disk groups by using ASMCMD and SQL*Plus.	None
Oracle Transparent Data Encryption (TDE)	Manage encryption wallets and encryption keys.	See SAP Notes <a href="#">2591575</a> and <a href="#">2799991</a> .
Oracle Database instances	Manage the SAP database and Oracle initialization parameters recommended by SAP.	For SAP-required Oracle initialization parameters, see SAP Notes <a href="#">2799900</a> , <a href="#">2470660</a> , and <a href="#">2470718</a> .
Oracle Recovery Manager (RMAN)	Back up, restore, and recover your SAP database.	None
Oracle Data Guard and Data Guard Broker	Configure and manage standby databases for disaster recovery scenarios.	See <a href="#">Introduction to Oracle Data Guard</a> .
\$IHRDBMS	The <i>installation</i> Oracle home of the database software.	For example: /u02/app/oracle/product/19.0.0.0/dbhome_1

ORACLE WORK AREA	PURPOSE	RELATED NOTES AND COMMENTS
\$OHRDBMS	The <i>runtime</i> Oracle home of the database software.	With SAP this is a symbolic link to the installation Oracle home of database software. For example: /oracle/MFG/19 -> /u02/app/oracle/product/19.0.0.0/dbhome_1
\$ORACLE_HOME	ORACLE_HOME environment variable	On the VM cluster nodes, it must be set to the value of \$OHRDBMS.

**Table 3: SAP-Specific Work Areas**

SAP WORK AREA	PURPOSE	RELATED NOTES AND COMMENTS
SAP Maintenance Planner	Create a <code>stack.xml</code> file for SAP Software Provisioning Manager (SWPM) and choose the SAP software components that you want to install.	See <a href="#">Maintenance Planner – User Guide</a> .
SAP Software Provisioning Manager (SWPM)	Use for <i>mandatory</i> tasks such as SWPM-based host preparation, installation of the SAP database instance and the primary application server (PAS), and for optional tasks such as the installation of ABAP system central services (ASCS), enqueue replication server (ERS), and additional application servers (AAS).	Always use the latest version of SWPM to avoid issues with new versions of Oracle Database software and new versions of Oracle Linux not supported in older versions of SWPM.
SAP NetWeaver software stack	Modify SAP instance profiles, configure RFC connections, and schedule database backups and checks in SAP transaction DB13.	SAP instance profiles must be adjusted to configure the correct number of work processes and SAP HA components if using SAPCTL.
SAProuter	Set up and configure SAProuter.	Customers must configure SAProuter at least for SAP EarlyWatch.
SAP GUI	Install SAP GUI components.	None
SAP HA using SAPCTL (optional)	Configure SAP HA for SAP ASCS and ERS by integrating with Oracle Clusterware.	See SAP Note <a href="#">1496927</a> to download the SAPCTL software package (includes documentation).
BR*Tools (optional)	Back up, restore, and recover your SAP database.	See SAP Notes <a href="#">1598594</a> , <a href="#">113747</a> , and <a href="#">776505</a> .

## Planning

Implementing Data Guard with SAP on Exadata Cloud Infrastructure or Oracle Exadata Cloud@Customer can require thorough planning, depending on customer requirements and the ability to manage potentially complex scenarios. The following sections touch on the key subject areas and must be elaborated in detail by customers to fit their specific scenarios and requirements.

### Scenarios and Complexity

Customers who are considering an implementation of Data Guard with SAP on Exadata Cloud Infrastructure or Oracle Exadata Cloud@Customer must answer an important question: What kind of disasters or scenarios do they want to handle and mitigate, and how much effort will it take to switch over from the primary site to the standby site?

## Delayed Standby

For some customers, it is important to simply have a standby database that reflects the current or a delayed state of the primary database. When incidents occur—like an administrative error—these customers can easily correct the mistake by either retrieving data from the standby database or switching over to it and then rebuilding the primary database from there, without ever actually running SAP productively against that standby.

## Read-Only Standby

For other customers, it might be important to use the standby for read-only purposes (without SAP NetWeaver Application Servers) or for running their backups.

## Full Disaster Recovery

Customers who cannot tolerate a prolonged total outage can turn the standby database into the new primary database and actually start SAP against it. Such a scenario arises when the data center running the SAP application and the Oracle Database becomes entirely unavailable (catastrophic outage).

However, this approach has significant challenges and difficulties. For example, the SAP database might have satellite systems, the SAP system might need RFC connections to other systems, or the system might use database links to reach out to other databases. Such dependencies make it challenging to bring up and properly run the SAP system on the standby site because the standby site might have different host names and IP addresses, the SAP enqueue lock table of the primary site will not be available, mount points might differ, SAP NetWeaver Application Server profiles could be different, SAP binaries must be kept in sync with the primary site, and many other factors. Customers must consider these issues and take appropriate measures to avoid them if they intend to activate the standby site and run SAP from there.

Some of the actions that are necessary to run SAP on the standby site must be done manually, while others can potentially be automated. However, the scope and the specific steps required are highly customer-specific and, therefore, cannot be described in this document. Customers themselves must identify the necessary measures and implement the appropriate actions so that their SAP system can run successfully in the standby environment.

We strongly recommend that customers thoroughly test, document, and communicate all measures necessary to switch the SAP system from the primary to the standby site, and to operate it on the standby site, so that in the event of a real incident, no issues or delays occur regarding these procedures.

The following diagram shows such a complex scenario with SAP High Availability implemented on the VM clusters using SAPCTL.

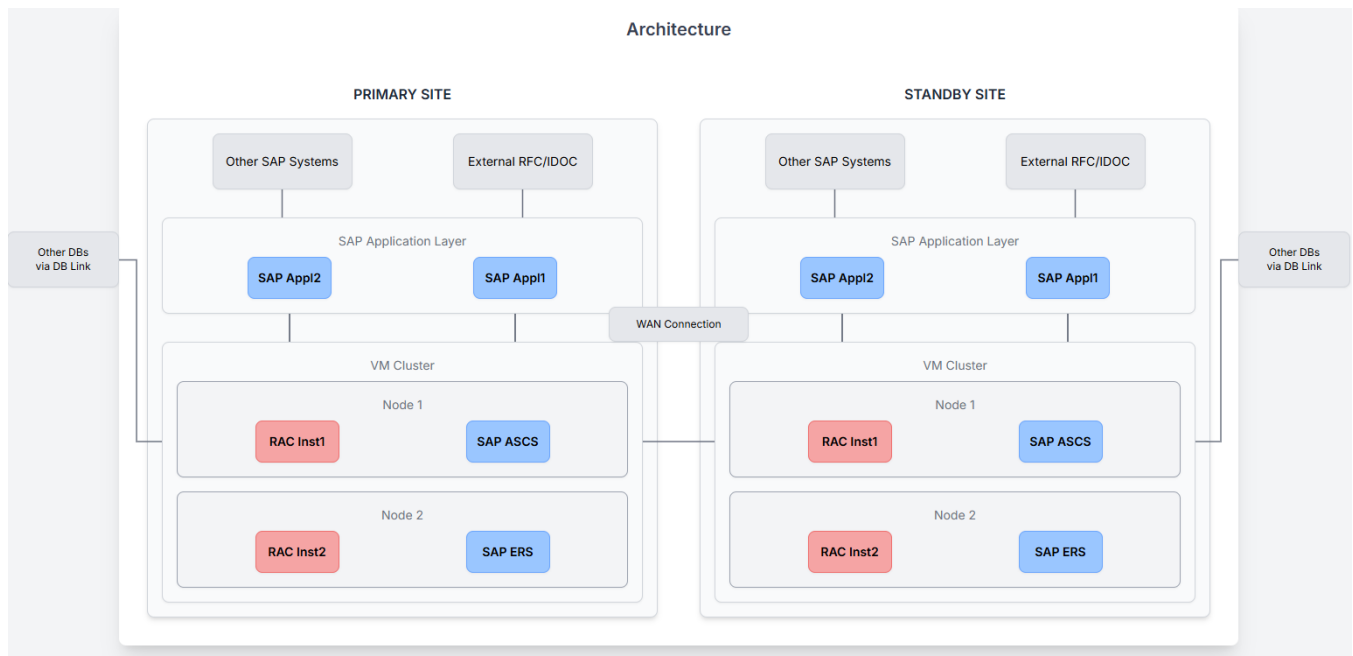


Figure 1: Architecture of a System with a Primary and Standby Site

**Note:** SAP NetWeaver Application Servers running on the primary site should never connect to the standby database, and servers running on the standby site should never connect to the primary database. Each SAP NetWeaver Application Server must use a `tnsnames.ora` entry that resolves only to the SCAN listener of the local site where the application server is deployed. The only exception would be an SAP instance started solely for maintenance purposes—for example, to perform configuration changes within the SAP system—and only if no productive operation is occurring at that time.

## Planning the Virtual Cloud Networks

Virtual cloud networks (VCNs) can be connected in several ways, including connecting a VCN to a customer-owned network. This document describes a setup where two VCNs are connected using dynamic routing gateways (DRGs)—one on the primary site and one on the standby site—within the same tenancy. It also explains how to configure an observer in a separate VCN as an optional component of a Data Guard configuration.

Customers who deploy the primary, standby, and, optionally, the observer VCN in different tenancies can also use this technique but must configure their tenancies to allow remote peering connections (RPCs) between trusted partners.

```
# In Tenancy A (the accepting side)
allow any-user from tenancy B_OCID to manage remote-peering-from in compartment
<compartment_name>
# In Tenancy B (the requesting side)
allow any-user from tenancy A_OCID to manage remote-peering-to in compartment <compartment_name>
```

Other ways to connect primary and standby VCNs—for example site-to-site VPN, FastConnect, or hub-and-spoke—are discussed in the [Oracle Cloud Infrastructure documentation](#).

The implementation of Oracle Data Guard across two locations requires a VCN on both the primary site and the standby site. Each VCN must contain at least two subnets: one for client access and one for backup operations. To ensure proper routing between the two client subnets, the respective CIDR blocks must not overlap.

For cross-compartment configurations, remote peering must be enabled using the required IAM policies. Following are examples of IAM policies for the compartments `cmp-london` and `cmp-frankfurt`:

```
allow group NetworkAdmins to manage remote-peering-from in compartment cmp-london
allow group NetworkAdmins to manage remote-peering-to in compartment cmp-frankfurt

allow group NetworkAdmins to manage remote-peering-from in compartment cmp-frankfurt
allow group NetworkAdmins to manage remote-peering-to in compartment cmp-london
```

In addition to well-configured network security groups (or security lists) and routing tables, both VCNs should also have a private view attached to the DNS resolver. This view should include a zone capable of resolving the SCAN name of the *remote* VM cluster.

A recommended approach is to first set up both VCNs (including the client and backup subnets), deploy both VM clusters as outlined in the next section, and then finalize and test the configuration. This approach includes the following steps:

1. Configure the DRGs.
2. Add the private view and DNS zone to the resolver.
3. Set up the necessary security lists and routing rules.
4. Verify communication between the primary and standby systems.

The following diagram illustrates an example VCN-to-VCN connection using DRGs.

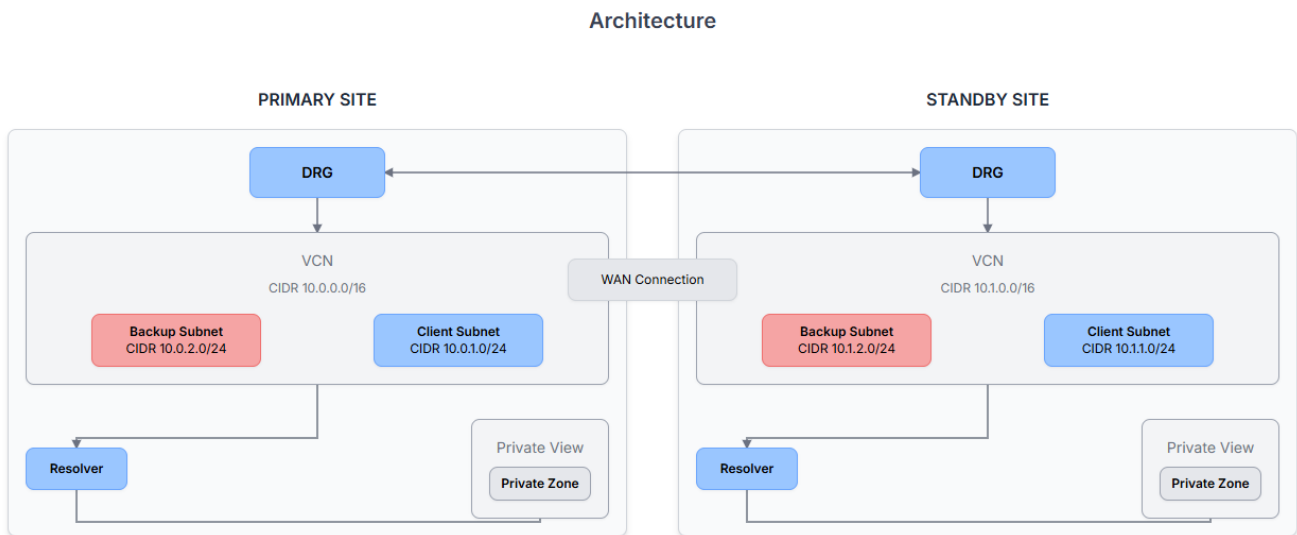


Figure 2: VCN-to-VCN Connection Using DRGs

## Planning the Deployment of the VM Clusters

VM clusters on both the primary site and the standby site must be deployed according to the SAP-specific technical briefs for the Oracle Exadata Database Service or Oracle Exadata Cloud@Customer platform, and must be configured. For consistent monitoring, troubleshooting, and correct interpretation of log and lag information, both the primary system and the standby system in a Data Guard setup must use the same OS and database time zone. Doing so avoids confusion with timestamps, prevents errors in SAP processes, and ensures reliable disaster recovery operations.

Setting up the standby database follows a procedure similar to migrating an SAP database to a VM cluster:

1. Prepare the VM cluster nodes for running SAP Software Provisioning Manager (SWPM).
2. Run the SAP SWPM host preparation on all VM cluster nodes.
3. Install the SAP database instance.
4. Remove the initially created SAP database from ASM, and either replace it with the database to be migrated or, in the case of setting up a standby database, copy the database from the primary site.

However, it should be noted that technical briefs for the respective platforms have changed over time, and not every Exadata-based platform is installed today in exactly the same way as in the past. Moreover, current technical briefs generally do not cover the additional use of Oracle Data Guard. Therefore, certain configuration adjustments are necessary after the initial installation of the primary and standby sites. This especially affects the location of the Oracle server parameter file (spfile), the TDE and autologin wallet, and the password file, which should all be located on ASM.

Before proceeding with any tasks on the primary and standby systems, it is essential to gather the names of the ASM disk groups as well as the hostnames of the SCAN listeners for both sites.

## Preparing the Primary Database

The steps required to prepare the primary database depend on whether the primary database was already configured with the initialization parameter `db_unique_name`. This parameter is required for Oracle Data Guard and also controls the subdirectory on ASM where the database stores datafiles, tempfiles, and controlfiles. You can set `db_unique_name` during the installation of a new SAP database instance in SWPM by adding it in the Oracle initialization parameters dialog where you set the `local_listener` parameters for the RAC instances. Note that not all locations of database related files are controlled by `db_unique_name`.

Customers who followed the corresponding technical brief for their platform should find the database-related files at the following locations.

**Table 4: Location of Database-Related Files**

FILE TYPE	LOCATION
DATAFILES, TEMPFILES db_unique_name not set db_unique_name set	+DATACn/[DB_NAME]/ +DATACn/[DB_UNIQUE_NAME]/
CONTROLFILES db_unique_name not set db_unique_name set	+DATACn/[DB_NAME]/ +RECOcN/[DB_NAME]/ +DATACn/[DB_UNIQUE_NAME]/ +RECOcN/[DB_UNIQUE_NAME]/
ONLINE REDO LOGFILES db_unique_name not set db_unique_name set	+RECOcN/[DB_NAME]/ +RECOcN/[DB_UNIQUE_NAME]/
ARCHIVELOGS	+RECOcN/[DB_NAME]/oraarch
TDE-WALLETS with wallet_root set to +DATACn/<DBNAME>/orawallet	+DATACn/[DB_NAME]/orawallet/tde  For example:  cwallet.sso => +DATAC4/MFG/AUTOLOGIN_KEY_STORE/cwallet.263.1205323817 ewallet.p12 => +DATAC4/MFG/KEY_STORE/ewallet.265.1205323817

FILE TYPE	LOCATION
Oracle server parameter file (spfile)	+DATAcN/[DB_NAME]/spfile<SID>.ora  For example:  spfilemfg.ora => +DATAc4/MFG/PARAMETERFILE/SPFILE.284.1205394607
Oracle initialization file (pfile)	\$ORACLE_HOME/dbs  For example:  initMFG.ora initMFG001.ora -> initMFG.ora
Oracle password file	\$ORACLE_HOME/dbs  For example:  orapwMFG orapwMFG001 -> orapwMFG

Customers who have not set db\_unique\_name on the primary database can choose one of the following options:

- Set db\_unique\_name to <DBSID> (for example, MFG) to keep all the files in their current location on ASM.
- Set db\_unique\_name to any other unique name (for example MFG\_PROD or MFG\_STBY) and move all the relevant files to the corresponding directories on ASM.

In the example described in this document, db\_unique\_name was set to MFG\_LON and all the files on ASM mentioned in the preceding table have been moved to subdirectory MFG\_LON (for London).

---

**Tip:** The most appropriate time to set or change db\_unique\_name is either during the migration to the primary system or when creating a new SAP database using SWPM, by setting the db\_unique\_name parameter when prompted for Oracle initialization parameters.

---

The example in this document configures the primary database in such a way that a standby database can be created from it. In the example we use MFG\_LON for the primary database (London) and MFG\_FFM for the standby database (Frankfurt).

- MFG\_LON runs on VM cluster nodes sc-xwyb81 and sc-xwyb82 and uses ASM disk groups +DATAc4 and +RECOc4.
- MFG\_FFM runs on VM cluster nodes sb-iwcan1 and sb-iwcan2 and uses ASM disk groups +DATAc8 and +RECOc8.

## Overview of Configuration Tasks on Primary Site

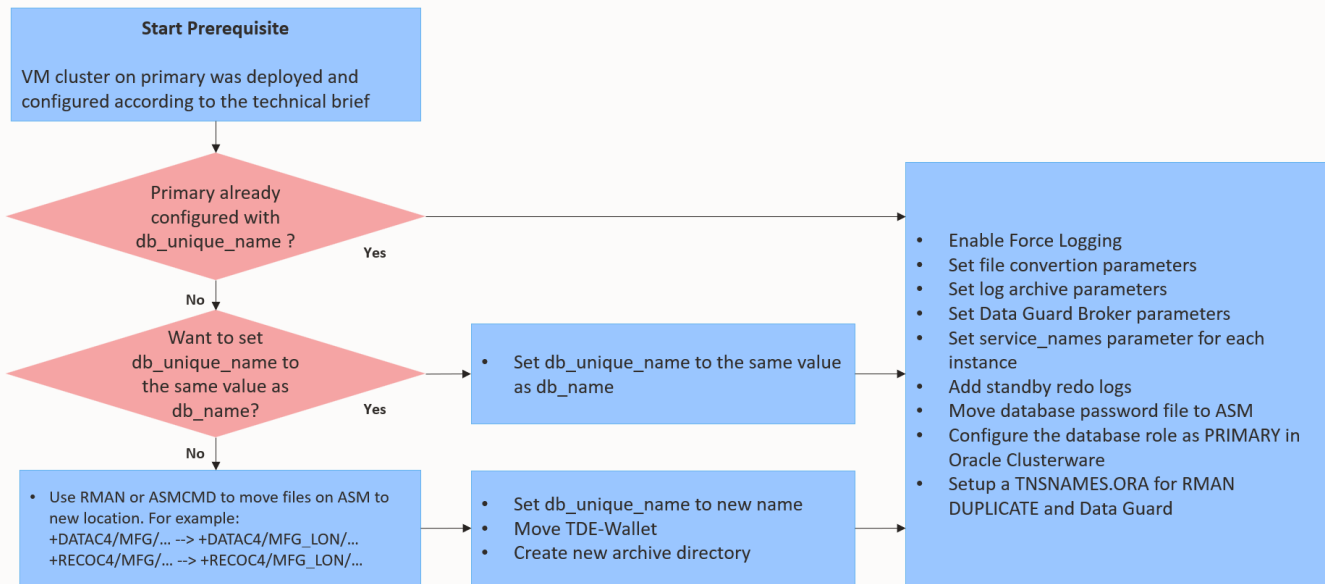


Figure 3: Configuration Tasks on the Primary Site

## Enable Force Logging on Primary Database

Run the following SQL command in sqlplus to enable force logging:

```
alter database force logging;
```

## Set File Name Conversion Parameters

Set the file name conversion parameters in case the primary database needs to assume the standby role at a later point in time.

```
alter system set db_file_name_convert='+DATA8','+DATA4' scope=spfile sid='*';
alter system set log_file_name_convert='+RECO8','+RECO4' scope=spfile sid='*';
```

## Set Log Archive Parameters

```
alter system set log_archive_config='dg_config=(MFG_LON,mfg_ffm)' scope=spfile sid='*';
alter system set log_archive_dest_1='LOCATION=+RECO4/MFG_LON/oraarch
VALID_FOR=(ALL_LOGFILES,ALL_ROLES) DB_UNIQUE_NAME=MFG_LON' scope=spfile sid='*';
```

## Set Oracle Data Guard Broker-Specific Parameters

```
alter system set dg_broker_config_file1='+DATA4/MFG_LON/dr1_mfg_lon.dat' scope=spfile sid='*';
alter system set dg_broker_config_file2='+DATA4/MFG_LON/dr2_mfg_lon.dat' scope=spfile sid='*';
alter system set dg_broker_start = TRUE scope=spfile sid='*';
```

## Set the SERVICE\_NAMES Parameter

**Warning:** Setting the SERVICE\_NAMES parameter with scope=both or scope=memory deregisters all services from Oracle TNS Listener started by Oracle Clusterware. As a result, no new connections to the database instances are possible. This may lead to unplanned downtime of SAP.

```
alter system set service_names='MFG_LON' scope=spfile sid='MFG001';
alter system set service_names='MFG_LON' scope=spfile sid='MFG002';
```

## Add Standby Redo Logs to the Primary Database

Add a complete set of standby redo logs for each thread, using the same size and number as the online redo logs. For example:

```
ALTER DATABASE ADD STANDBY LOGFILE THREAD 1
  GROUP 100 ('+RECO4') SIZE 200M,
  GROUP 101 ('+RECO4') SIZE 200M,
  GROUP 102 ('+RECO4') SIZE 200M,
  GROUP 103 ('+RECO4') SIZE 200M;
```

```
ALTER DATABASE ADD STANDBY LOGFILE THREAD 2
  GROUP 200 ('+RECO4') SIZE 200M,
  GROUP 201 ('+RECO4') SIZE 200M,
  GROUP 202 ('+RECO4') SIZE 200M,
  GROUP 203 ('+RECO4') SIZE 200M;
```

## Move the Database Password File to ASM

As mentioned earlier, the database password file is located at \$ORACLE\_HOME/dbs after initial installation. As long as the database password file is located on the file system, changing the password (for example of user SYS) requires changing it on each RAC instance. To allow password changes of a database user in the password file to be shared across all RAC instances, move the database password file to ASM and remove the password file and the corresponding symbolic link under \$ORACLE\_HOME/dbs.

Copying the database password file with the --dbuniquefilename option also updates the corresponding database resource in Oracle Clusterware.

```
[oracle@sc-xwyb81 ~]$ cd $ORACLE_HOME/dbs
[oracle@sc-xwyb81 dbs]$ ls -la orapw*
-rw-r----- 1 oracle oinstall 6144 Jul 23 08:26 orapwMFG
lrwxrwxrwx 1 oracle oinstall    8 Jul 23 07:32 orapwMFG001 -> orapwMFG
[oracle@sc-xwyb81 dbs]$ pwd
/oracle/MFG/19/dbs
[oracle@sc-xwyb81 dbs]$ asmcmd
ASMCMD> pwcopy --dbuniquefilename MFG_LON /oracle/MFG/19/dbs/orapwMFG +DATAC4/MFG_LON/orapwMFG
copying /oracle/MFG/19/dbs/orapwMFG -> +DATAC4/MFG_LON/orapwMFG
ASMCMD> exit
[oracle@sc-xwyb81 dbs]$ rm orapwMFG
[oracle@sc-xwyb81 dbs]$ rm orapwMFG001
```

Remove the database password file and symbolic link on the remaining VM cluster nodes

```
[oracle@sc-xwyb8N dbs]$ rm orapwMFG
[oracle@sc-xwyb8N dbs]$ rm orapwMFG00N
```

## Move the TDE Wallet

If you set `db_unique_name` to a value different than `db_name`—for example, `db_unique_name` is now `MFG_LON`, whereas it was different or unset previously—you must move your TDE wallet and autologin wallet to the new directory—for example, from `+DGNAME/MFG/orawallet` to `+DGNAME/MFG_LON/orawallet`.

Ensure that the directory structure on ASM exists.

```
ASMCMD> mkdir +DATAC4/MFG_LON
ASMCMD> mkdir +DATAC4/MFG_LON/orawallet
ASMCMD> mkdir +DATAC4/MFG_LON/orawallet/tde
```

Copy the TDE wallet and autologin wallet to the new location.

```
ASMCMD> cp --dest_dbname MFG_LON +DATAC4/MFG/orawallet/tde/cwallet.sso
+DATAC4/MFG_LON/orawallet/tde
copying +DATAC4/MFG/orawallet/tde/cwallet.sso -> +DATAC4/MFG_LON/orawallet/tde/cwallet.sso
ASMCMD> cp --dest_dbname MFG_LON +DATAC4/MFG/orawallet/tde/ewallet.p12
+DATAC4/MFG_LON/orawallet/tde
copying +DATAC4/MFG/orawallet/tde/ewallet.p12 -> +DATAC4/MFG_LON/orawallet/tde/ewallet.p12

ASMCMD> ls -la +DATAC4/MFG_LON/orawallet/tde/*
Type      Redund  Striped  Time          Sys  Name
KEY_STORE HIGH    COARSE   JUL 23 11:00:00 N    cwallet.sso =>
+DATAC4/MFG_LON/KEY_STORE/cwallet.sso.301.1207221285
KEY_STORE HIGH    COARSE   JUL 23 11:00:00 N    ewallet.p12 =>
+DATAC4/MFG_LON/KEY_STORE/ewallet.p12.300.1207221383
```

## Create New Archive Directory

If you set `db_unique_name` to a value different from `db_name`—for example, `db_unique_name` is now `MFG_LON`, whereas it was different or unset previously—you must create a new `oraarch` directory.

```
ASMCMD> mkdir +RECOC4/MFG_LON
ASMCMD> mkdir +RECOC4/MFG_LON/oraarch
```

## Configure Database as Primary in Oracle Clusterware

```
srvctl modify database -db MFG_LON -role PRIMARY
```

## Restart Database Instances on the Primary Using srvctl

You must restart the instances for the configuration changes to take effect.

```
srvctl stop database -db MFG_LON
srvctl start database -db MFG_LON
```

## Set Up TNSNAMES.ORA for RMAN DUPLICATE

To allow the primary database to connect to the standby database, the corresponding entries must be added to `$ORACLE_HOME/network/admin/tnsnames.ora`. As shown in the following example, the entry `MFG_LON` points to the SCAN listener of the primary database, and `MFG_FFM` points to the SCAN listener of the standby database. As mentioned earlier, the corresponding DNS record (the SCAN name with the three SCAN VIPs of the standby) must exist.

Additionally, the example shows using a temporary listener on the standby site to set up the standby database. The MFG\_DUP entry points to one of the node IP addresses on the standby site where this listener will be started. The advantage of this approach is that customers do not need to modify the `listener.ora` file in the Grid Infrastructure home, which is not intended to be modified by customers.

```
MFG_FFM =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = <fqdn_scanname_ffm>)(PORT = 1521))
    (CONNECT_DATA = (SERVICE_NAME = MFG_FFM))
  )

MFG_LON =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = <fqdn_scanname_lon>)(PORT = 1521))
    (CONNECT_DATA = (SERVICE_NAME = MFG_LON))
  )

MFG_DUP =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = <nodeip_ffm>)(PORT = 1621))
    (CONNECT_DATA = (SERVICE_NAME = MFG001))
  )
```

## Final Checks on the Primary

If you set the `db_unique_name` to a value different from `db_name`—for example, `db_unique_name` is now `MFG_LON`, whereas it was different or unset previously—you should check all ASM disk groups to ensure they no longer contain directories or files under the old name.

For example, if your database was originally located under `+DGNAME/MFG` and is now located under `+DGNAME/MFG_LON`, the directory `+DGNAME/MFG` should no longer exist.

If the directory still exists, check whether any references point to a location outside of `MFG`, and remove any files that are no longer required.

## Setting Up the Standby Site

The following sections describe the steps to set up the standby site. The tasks assume that customers ran SWPM host preparation and installed the SAP database instance according to the technical brief for their platform.

## Overview of Configuration Tasks on Standby Site

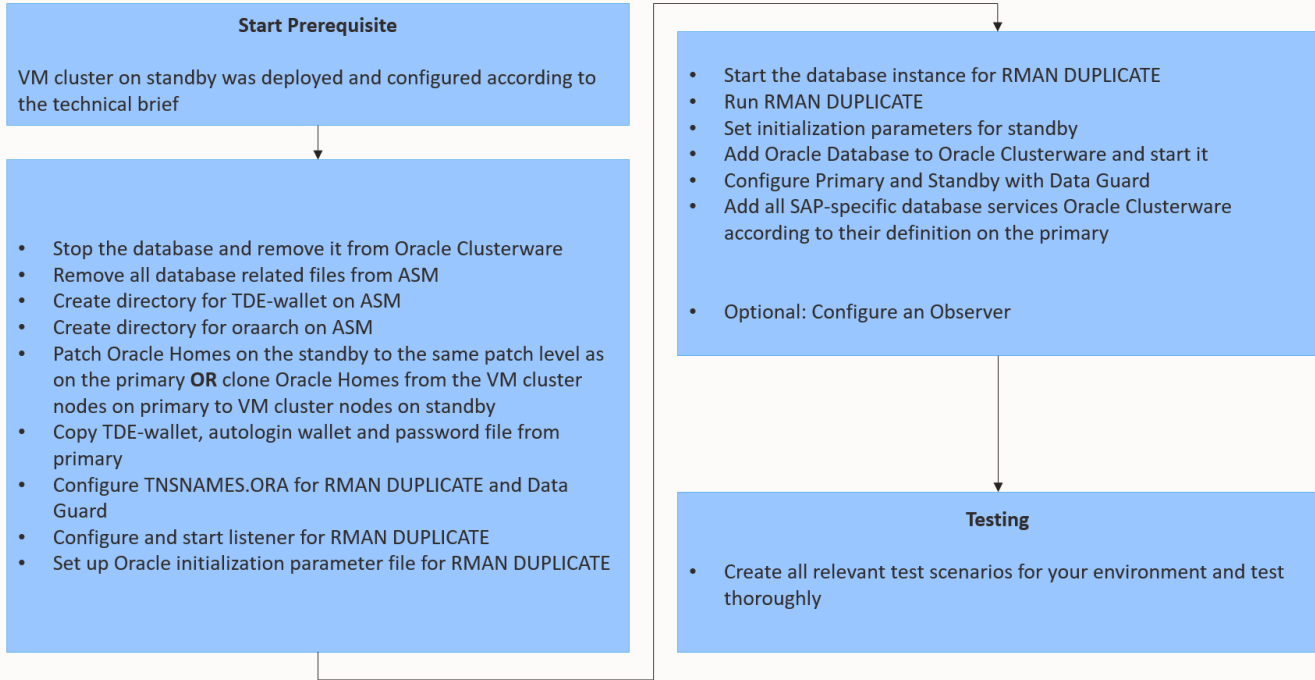


Figure 4: Configuration Tasks on the Secondary Site

## Stop the SAP Database on Standby and Remove It from Oracle Clusterware

The SWPM-based SAP database instance installation is required to install and configure the SAP environment on the VM cluster nodes. The SAP database that is created during that process is removed from ASM and Oracle Clusterware because it is completely replaced when the standby database is copied from the primary.

```

[root@sb-iwcan1 ~]# su - oracle
[oracle@sb-iwcan1 ~]$ . oraenv
ORACLE_SID = [oracle] ? MFG
The Oracle base has been set to /u02/app/oracle
[oracle@sb-iwcan1 ~]$ srvctl stop database -db MFG
[oracle@sb-iwcan1 ~]$ srvctl remove database -db MFG
  
```

## Remove the SAP Database from ASM

```

[oracle@sb-iwcan1 ~]$ asmcmd
ASMCMD> cd DATA8
ASMCMD> ls
ASM/
MFG/
c1-6ggfxa6q/
orapwasm
orapwasm_backup
ASMCMD> rm -rf MFG
ASMCMD> cd ..
ASMCMD> cd RECOC8
ASMCMD> ls
MFG /
c1-6ggfxa6q/
ASMCMD> rm -rf MFG
  
```

## Create TDE Wallet and Archive Directories on ASM

```
[oracle@sb-iwcan1 ~]$ asmcmd
ASMCMDB> mkdir +DATAC8/MFG_FFM
ASMCMDB> mkdir +DATAC8/MFG_FFM/orawallet
ASMCMDB> mkdir +DATAC8/MFG_FFM/orawallet/tde
ASMCMDB> mkdir +RECOC8/MFG_FFM
ASMCMDB> mkdir +RECOC8/MFG_FFM/oraarch
```

## Copy the TDE Wallet, the Autologin Wallet, and the Password File from Primary

On the primary site, copy the TDE wallet, the autologin wallet, and the password file to the local file system. For example:

```
ASMCMDB> cd +DATAC4/MFG_LON/orawallet/tde
ASMCMDB> cp ewallet.p12 /tmp/mfglon_ewallet.p12
copying +DATAC4/MFG_LON/orawallet/tde/ewallet.p12 -> /tmp/mfglon_ewallet.p12
ASMCMDB> cp cwallet.sso /tmp/mfglon_cwallet.sso
copying +DATAC4/MFG_LON/orawallet/tde/cwallet.sso -> /tmp/mfglon_cwallet.sso
ASMCMDB> cp +DATAC4/MFG_LON/orapwmfglon /tmp/orapwmfglon
copying +DATAC4/MFG_LON/orapwmfglon -> /tmp/orapwmfglon
```

Copy the files from the primary site to the standby site. For example:

```
[opc@sc-xwyb81 ~]$ sudo chmod 744 /tmp/mfglon_*
[opc@sc-xwyb81 ~]$ sudo chmod 744 /tmp/orapwmfglon
[opc@sc-xwyb81 ~]$ scp -i myssh.key /tmp/mfglon_* <node1-ip-ffm>:/tmp
[opc@sc-xwyb81 ~]$ scp -i myssh.key /tmp/orapwmfglon <node1-ip-ffm>:/tmp
```

On the standby site, copy the files to ASM. For example:

```
ASMCMDB> cp --dest_dbname MFG_FFM /tmp/mfglon_ewallet.p12
+DATAC8/MFG_FFM/orawallet/tde/ewallet.p12
copying /tmp/mfglon_ewallet.p12 -> +DATAC8/MFG_FFM/orawallet/tde/ewallet.p12
ASMCMDB> cp --dest_dbname MFG_FFM /tmp/mfglon_cwallet.sso
+DATAC8/MFG_FFM/orawallet/tde/cwallet.sso
copying /tmp/mfglon_cwallet.sso -> +DATAC8/MFG_FFM/orawallet/tde/cwallet.sso

ASMCMDB> pwcopyp --dbuniquename MFG_FFM /tmp/orapwmfglon +DATAC8/MFG_FFM/orapwmfg_ffm -f
copying /tmp/orapwmfglon -> +DATAC8/MFG_FFM/orapwmfg_ffm
```

On the standby site, copy the password file to \$ORACLE\_HOME/dbs/orapwMFG001. This is required only for running RMAN DUPLICATE and only on one VM cluster node. The password file must be deleted after setting up the standby database.

```
oracle@sb-iwcan1 dbs]$ cp /tmp/orapwmfglon $ORACLE_HOME/dbs/orapwMFG001
```

## Configure TNSNAMES.ORA

Add the entries shown in the earlier section, “Set Up TNSNAMES.ORA for RMAN DUPLICATE.”

## Configure and Start the Listener for RMAN DUPLICATE

On the standby site, create a `listener.ora` file in `$ORACLE_HOME/network/admin`. The following example uses port 1621 because 1521 is reserved for the listener controlled by Grid Infrastructure.

```
[oracle@sb-iwcan1 admin]$ cat listener.ora
LISTENER_DUP =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS=(PROTOCOL=TCP)(HOST=<node1-ip-ffm>)(PORT=1621))
    )
  )

SID_LIST_LISTENER_DUP =
  (SID_LIST =
    (SID_DESC =
      (SID_NAME = MFG001)
      (ORACLE_HOME = /oracle/MFG/19)
    )
  )
```

Start the listener.

```
[oracle@sb-iwcan1 admin]$ lsnrctl start LISTENER_DUP
LSNRCTL for Linux: Version 19.0.0.0.0 - Production on 24-JUL-2025 14:31:54
Copyright (c) 1991, 2025, Oracle. All rights reserved.
Starting /oracle/MFG/19/bin/tnslsnr: please wait...
TNSLSNR for Linux: Version 19.0.0.0.0 - Production
System parameter file is /oracle/MFG/19/network/admin/listener.ora
Log messages written to /u02/app/oracle/diag/tnslsnr/sb-iwcan1/listener_dup/alert/log.xml
Listening on: (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=<node1-ip-ffm>)(PORT=1621)))
Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=TCP)(HOST=<node1-ip-ffm>)(PORT=1621)))
STATUS of the LISTENER
-----
Alias                LISTENER_DUP
Version              TNSLSNR for Linux: Version 19.0.0.0.0 - Production
Start Date           24-JUL-2025 14:31:54
Uptime               0 days 0 hr. 0 min. 0 sec
Trace Level          off
Security             ON: Local OS Authentication
SNMP                 OFF
Listener Parameter File /oracle/MFG/19/network/admin/listener.ora
Listener Log File    /u02/app/oracle/diag/tnslsnr/sb-iwcan1/listener_dup/alert/log.xml
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=<node1-ip-ffm>)(PORT=1621)))
Services Summary...
Service "MFG001" has 1 instance(s).
  Instance "MFG001", status UNKNOWN, has 1 handler(s) for this service...
The command completed successfully
```

## Create an Oracle Initialization Parameter File for RMAN DUPLICATE

On the standby site, create an Oracle initialization parameter file to start the initial database instance into the nomount state for RMAN DUPLICATE to build the standby database.

```
[oracle@sb-iwcan1 dbs]$ cat init_for_dup.ora
*.db_name='MFG'
*.db_create_file_dest='+DATAC8'
*.db_create_online_log_dest_1='+RECO8'
*.db_recovery_file_dest='+RECO8'
*.db_recovery_file_dest_size=200G
*.db_unique_name='MFG_FFM'
*.db_file_name_convert='+DATAC4','+DATAC8'
*.log_file_name_convert='+RECO4','+RECO8'
*.diagnostic_dest='/oracle/MFG/saptrace'
MFG001.instance_name='MFG001'
MFG001.instance_number=001
MFG001.local_listener='<node1-ip-ffm>:1621'
*.log_archive_config='DG_CONFIG=(MFG_LON,MFG_FFM)'
*.log_archive_dest_1='LOCATION=+RECO8/MFG_FFM/oraarch VALID_FOR=(ALL_LOGFILES,ALL_ROLES)
DB_UNIQUE_NAME='MFG_FFM'
*.log_archive_dest_2='SERVICE=MFG_LON ASYNC NOAFFIRM VALID_FOR=(ONLINE_LOGFILES,PRIMARY_ROLE)
DB_UNIQUE_NAME=MFG_LON'
*.log_archive_format='%t_%s_%r.dbf'
*.remote_login_passwordfile='exclusive'
MFG001.service_names='MFG_FFM'
*.tde_configuration='KEYSTORE_CONFIGURATION=FILE'
MFG001.thread=001
MFG001.undo_tablespace='PSAPUNDO001'
*.wallet_root='+DATAC8/MFG_FFM/orawallet'
```

## Start the Database Instance for RMAN DUPLICATE

```
[root@sb-iwcan1 ~]# su - oracle
[oracle@sb-iwcan1 ~]$ . oraenv
ORACLE_SID = [oracle] ? MFG
The Oracle base has been set to /u02/app/oracle
[oracle@sb-iwcan1 ~]$ export ORACLE_SID=MFG001
[oracle@sb-iwcan1 ~]$ cd $ORACLE_HOME/dbs
[oracle@sb-iwcan1 dbs]$ sqlplus / as sysdba

SQL*Plus: Release 19.0.0.0.0 - Production on Fri Jul 25 13:50:08 2025
Version 19.27.0.0.0

Copyright (c) 1982, 2024, Oracle. All rights reserved.

Connected to an idle instance.

SQL> startup nomount pfile=init_for_dup.ora
ORACLE instance started.

Total System Global Area  365819896 bytes
Fixed Size                  8939512 bytes
Variable Size              281018368 bytes
```

```

Database Buffers          67108864 bytes
Redo Buffers              8753152 bytes
SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.27.0.0.0

```

## Run RMAN DUPLICATE

The following example uses 8 RMAN channels in parallel to speed up duplication.

```

rman target sys/123_BETA@MFG_LON auxiliary sys/123_BETA@MFG_DUP

run {
ALLOCATE CHANNEL t1 DEVICE TYPE disk;
ALLOCATE CHANNEL t2 DEVICE TYPE disk;
ALLOCATE CHANNEL t3 DEVICE TYPE disk;
ALLOCATE CHANNEL t4 DEVICE TYPE disk;
ALLOCATE CHANNEL t5 DEVICE TYPE disk;
ALLOCATE CHANNEL t6 DEVICE TYPE disk;
ALLOCATE CHANNEL t7 DEVICE TYPE disk;
ALLOCATE CHANNEL t8 DEVICE TYPE disk;
ALLOCATE AUXILIARY CHANNEL a1 DEVICE TYPE disk;
ALLOCATE AUXILIARY CHANNEL a2 DEVICE TYPE disk;
ALLOCATE AUXILIARY CHANNEL a3 DEVICE TYPE disk;
ALLOCATE AUXILIARY CHANNEL a4 DEVICE TYPE disk;
ALLOCATE AUXILIARY CHANNEL a5 DEVICE TYPE disk;
ALLOCATE AUXILIARY CHANNEL a6 DEVICE TYPE disk;
ALLOCATE AUXILIARY CHANNEL a7 DEVICE TYPE disk;
ALLOCATE AUXILIARY CHANNEL a8 DEVICE TYPE disk;

duplicate target database
  for standby
  from active database
  dorecover
  spfile
    set CONTROL_FILES='+DATAC8/MFG_FFM/cntrlmfg.dbf','+RECOC8/MFG_FFM/cntrlmfg.dbf'
    set DB_UNIQUE_NAME='MFG_FFM'
    set DB_CREATE_FILE_DEST='+DATAC8'
    set DB_CREATE_ONLINE_LOG_DEST_1='+RECOC8'
    set log_archive_dest_1='LOCATION=+RECOC8/MFG_FFM/oraarch VALID_FOR=(ALL_LOGFILES,ALL_ROLES)
    DB_UNIQUE_NAME='MFG_FFM'
    set FAL_SERVER='MFG_LON'
    set FAL_CLIENT='MFG_FFM'
    set STANDBY_FILE_MANAGEMENT='AUTO'
    set DG_BROKER_START='TRUE'
    set DB_FILE_NAME_CONVERT='+DATAC4','+DATAC8'
    set LOG_FILE_NAME_CONVERT='+RECOC4','+RECOC8'
    set wallet_root='+DATAC8/MFG_FFM/orawallet'
    set db_recovery_file_dest='+RECOC8'
  nofilenamecheck;
}

```

## Set Initialization Parameters for Standby

```
[oracle@sb-iwcan1 dbs]$ sqlplus / as sysdba
SQL*Plus: Release 19.0.0.0.0 - Production on Fri Jul 25 14:37:51 2025
Version 19.27.0.0.0
Copyright (c) 1982, 2024, Oracle. All rights reserved.
Connected to:
Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.27.0.0.0

alter system set local_listener='sb-iwcan1-vip:1521' scope=both sid='MFG001';
alter system set local_listener='sb-iwcan2-vip:1521' scope=both sid='MFG002';
alter system set remote_listener='sb-iwcan-scan.sapexacert2clie.sapexacert2vcn.oraclevcn.com'
scope=both sid='*';
alter system set service_names='MFG_FFM' scope=both sid='MFG001';
alter system set service_names='MFG_FFM' scope=both sid='MFG002';
alter system set db_recovery_file_dest_size=300G scope=both sid='*';
alter system set dg_broker_config_file1 = '+DATAC8/MFG_FFM/DR1_MFG_FFM.DAT' scope=spfile
sid='*';
alter system set dg_broker_config_file2 = '+DATAC8/MFG_FFM/DR2_MFG_FFM.DAT' scope=spfile
sid='*';
alter system set dg_broker_start = TRUE scope=both sid='*';
alter system reset log_archive_dest_2 scope=both sid='*';
```

Shut down the database and remove the password file \$ORACLE\_HOME/dbs/initMFG001.ora.

```
SQL> shutdown immediate;
Database dismounted.
ORACLE instance shut down.
SQL> exit
Disconnected from Oracle Database 19c EE Extreme Perf Release 19.0.0.0.0 - Production
Version 19.27.0.0.0
[oracle@sb-iwcan1 dbs]$ rm orapwMFG001
```

## Configure the Database in Clusterware and Start It

For example:

```
oracle@sb-iwcan1 dbs]$ srvctl add database -db MFG_FFM -dbname MFG -oraclehome /oracle/MFG/19 -
spfile +DATAC8/MFG_FFM/spfileMFG.ora -pwfile +DATAC8/MFG_FFM/orapwmfg_ffm -role PHYSICAL_STANDBY
[oracle@sb-iwcan1 dbs]$ srvctl add instance -db MFG_FFM -instance MFG001 -node sb-iwcan1
[oracle@sb-iwcan1 dbs]$ srvctl add instance -db MFG_FFM -instance MFG002 -node sb-iwcan2
[oracle@sb-iwcan1 dbs]$ srvctl start database -db MFG_FFM
```

## Configure the Primary and Standby with Data Guard

```
DGMGRL> CREATE CONFIGURATION DG_MFG_LON_FFM AS PRIMARY DATABASE IS MFG_LON CONNECT IDENTIFIER IS
MFG_LON;
Configuration "dg_mfg_lon_ffm" created with primary database "mfg_lon"
DGMGRL> ADD DATABASE MFG_FFM AS CONNECT IDENTIFIER IS MFG_FFM MAINTAINED AS PHYSICAL;
Database "mfg_ffm" added
DGMGRL> enable configuration;
Enabled.
DGMGRL> show configuration;
```

## Configuration - dg\_mfg\_lon\_ffm

```
Protection Mode: MaxPerformance
Members:
mfg_lon - Primary database
mfg_ffm - Physical standby database
```

Fast-Start Failover: Disabled

```
Configuration Status:
SUCCESS (status updated 9 seconds ago)
```

```
DGMGRL> edit database MFG_LON SET PROPERTY LogXptMode=SYNC;
DGMGRL> Property "logxptmode" updated
DGMGRL> edit database MFG_FFM SET PROPERTY LogXptMode=SYNC;
DGMGRL> Property "logxptmode" updated
DGMGRL> edit database MFG_FFM SET STATE='APPLY-ON';
Succeeded.
DGMGRL> show configuration
```

## Configuration - dg\_mfg\_lon\_ffm

```
Protection Mode: MaxPerformance
Members:
mfg_lon - Primary database
mfg_ffm - Physical standby database
```

Fast-Start Failover: Disabled

```
Configuration Status:
SUCCESS (status updated 58 seconds ago)
DGMGRL> EDIT CONFIGURATION SET PROTECTION MODE AS MAXAVAILABILITY;
Succeeded.
```

## Add SAP-Specific Database Services to Oracle Clusterware

Add all relevant SAP-specific database services to the database in Oracle Clusterware. Check the configuration of each of those services on the primary and add it on the standby with the exact same configuration parameters.

For example, on the primary, run the following commands:

```
[oracle@sc-xwyb81 ~]$ srvctl config service -db MFG_LON -s MFG_D10
Service name: MFG_D10
Server pool:
Cardinality: 1
Service role: PRIMARY
Management policy: AUTOMATIC
DTP transaction: false
AQ HA notifications: true
Global: false
Commit Outcome: false
Failover type:
Failover method:
Failover retries:
```

```

Failover delay:
Failover restore: NONE
Connection Load Balancing Goal: LONG
Runtime Load Balancing Goal: NONE
TAF policy specification: NONE
Edition:
Pluggable database name:
Hub service:
Maximum lag time: ANY
SQL Translation Profile:
Retention: 86400 seconds
Replay Initiation Time: 300 seconds
Drain timeout:
Stop option:
Session State Consistency: DYNAMIC
GSM Flags: 0
Service is enabled
Preferred instances: MFG001
Available instances: MFG002
CSS critical: no

```

Then, create the same service on the standby, for example:

```

[oracle@sb-iwcan1 ~]$ srvctl add service -db MFG_FFM -service MFG_D10 -notification TRUE -
preferred MFG001 -available MFG002
[oracle@sb-iwcan1 ~]$ srvctl config service -db MFG_FFM -service MFG_D10
Service name: MFG_D10
Server pool:
Cardinality: 1
Service role: PRIMARY
Management policy: AUTOMATIC
DTP transaction: false
AQ HA notifications: true
Global: false
Commit Outcome: false
Failover type:
Failover method:
Failover retries:
Failover delay:
Failover restore: NONE
Connection Load Balancing Goal: LONG
Runtime Load Balancing Goal: NONE
TAF policy specification: NONE
Edition:
Pluggable database name:
Hub service:
Maximum lag time: ANY
SQL Translation Profile:
Retention: 86400 seconds
Replay Initiation Time: 300 seconds
Drain timeout:
Stop option:
Session State Consistency: DYNAMIC
GSM Flags: 0

```

```
Service is enabled
Preferred instances: MFG001
Available instances: MFG002
CSS critical: no
```

## Implementing the Fast Start Failover (Optional)

Implementing Fast-Start Failover requires the setup of a third site where the Data Guard Observer runs. Customers who understand the implications of an automated failover in the event of a primary site outage—particularly in the context of SAP systems—can use the Data Guard Observer to monitor their primary site and automatically promote the standby database if needed.

However, it is the customer’s responsibility to ensure that all necessary measures are in place to successfully run SAP on the standby site after a failover. Simply activating the standby database is not sufficient in most cases, as nearly all SAP systems have dependencies that extend beyond database availability. These dependencies may include connected satellite systems, RFC destinations, database links, and other integrations.

Customers who cannot operate their SAP system on the standby site without additional manual intervention should carefully consider whether to enable Fast-Start Failover at all—or whether it is preferable to perform a manual switchover based on their own assessment of the situation.

To deploy a Data Guard Observer in a third site, an *additional virtual cloud network (VCN)* is required and *IAM policies* must be adjusted for cross-compartment or cross-tenancy remote connection peering. The example in this document shows this VCN connected to both the primary site VCN and the standby site VCN using dynamic routing gateways (DRGs) and remote peering connections (RPCs). In addition, appropriate private DNS zones and views must be created to ensure that the DNS resolver can resolve the SCAN IP addresses of both the primary and standby sites. Proper routing rules must also be configured to ensure that traffic is forwarded to the correct network segment in each case.

---

**Important:** For Fast-Start Failover, customers must enable flashback database on the primary and standby sites.

---

The following diagram illustrates this setup.

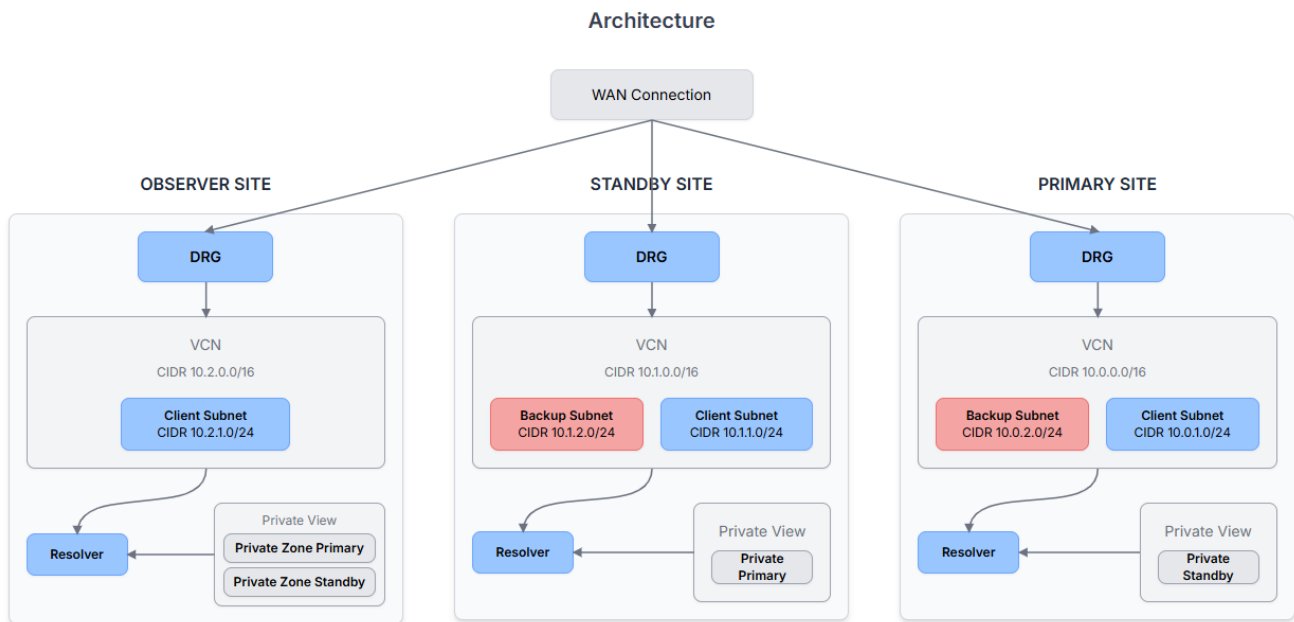


Figure 5: Configuration of the Data Guard Observer

To configure a Data Guard Observer, you must deploy a compute instance in which the Observer can run. The easiest way to do this is to install a full Oracle Database Home, for example by downloading and installing a preconfigured RPM from Oracle.

For example:

1. Download `oracle-database-ee-19c-1.0-1.x86_64.rpm` from Oracle.

2. Install the RPM.

```
dnf -y install oracle-database-ee-19c-1.0-1.x86_64.rpm
```

3. Add the following lines to `.bash_profile` of the `oracle` user.

```
# User specific environment and startup programs
export ORACLE_BASE=/opt/oracle
export ORACLE_HOME=/opt/oracle/product/19c/dbhome_1
export LD_LIBRARY_PATH=$ORACLE_HOME/lib
export PATH=$PATH:$ORACLE_HOME/bin
export TNS_ADMIN=$ORACLE_HOME/network/admin
```

4. As the `oracle` user, relink all the Oracle binaries.

```
[oracle@observer ~]$ cd $ORACLE_HOME/bin
[oracle@observer bin]$ ./relink all
```

5. Set up `tnsnames.ora`.

```
oracle@observer ~]$ cat $ORACLE_HOME/network/admin/tnsnames.ora
MFG_FFM =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = sb-iwcan-
scan.sapexacert2clie.sapexacert2vcn.oraclevcn.com)(PORT = 1521))
    (CONNECT_DATA = (SERVICE_NAME = MFG_FFM))
  )

MFG_LON =
  (DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP)(HOST = sc-xwyb8-
scan.sapexacert1clie.sapexacert1vcn.oraclevcn.com)(PORT = 1521))
    (CONNECT_DATA = (SERVICE_NAME = MFG_LON))
  )
```

6. Check network connectivity.

```
[oracle@observer ~]$ tnsping MFG_LON
[oracle@observer ~]$ tnsping MFG_FFM
[oracle@observer ~]$ sqlplus SYS/123_BETA@MFG_LON AS SYSDBA
[oracle@observer ~]$ sqlplus SYS/123_BETA@MFG_FFM AS SYSDBA
```

7. Use `dgmgrl` to enable Fast-Start Failover.

```
oracle@observer ~]$ dgmgrl SYS/123_BETA@MFG_LON AS SYSDBA
DGMGRL for Linux: Release 19.0.0.0.0 - Production on Tue Jul 29 09:49:37 2025
Version 19.3.0.0.0

Copyright (c) 1982, 2019, Oracle and/or its affiliates. All rights reserved.

Welcome to DGMGRL, type "help" for information.
Connected to "MFG_LON"
```

```
Connected as SYSDBA.
DGMGRL> enable fast_start failover;
```

8. Configure and start the observer.

```
EDIT CONFIGURATION SET PROPERTY FastStartFailoverThreshold = 30;
EDIT CONFIGURATION SET PROPERTY FastStartFailoverLagLimit = 30;
EDIT CONFIGURATION SET PROPERTY FastStartFailoverTarget = 'MFG_FFM';
START OBSERVER FILE='/home/oracle/observer.log';
```

## Testing the Configuration

Customers should carefully consider potential failure scenarios, their impact, and the necessary actions well in advance. We strongly recommend walking through all relevant scenarios in detail to ensure that the planned measures are effective in a real outage situation and to avoid complications, for example, caused by human error.

The switchover outlined in the following example should be understood as a basic functional test only.

```
DGMGRL> switchover to 'MFG_FFM';
Performing switchover NOW, please wait...
Operation requires a connection to database "mfg_ffm"
Connecting ...
Connected to "MFG_FFM"
Connected as SYSDBA.
New primary database "mfg_ffm" is opening...
Oracle Clusterware is restarting database "mfg_lon" ...
Connected to "MFG_LON"
Switchover succeeded, new primary is "MFG_FFM"
DGMGRL> switchover to 'MFG_LON';
Performing switchover NOW, please wait...
Operation requires a connection to database "mfg_lon"
Connecting ...
Connected to "MFG_LON"
Connected as SYSDBA.
New primary database "mfg_lon" is opening...
Oracle Clusterware is restarting database "mfg_ffm" ...
Connected to "MFG_FFM"
Switchover succeeded, new primary is "MFG_LON"
```

## Installing SAP NetWeaver Application Server ABAP/Java

Customers who have decided to set `db_unique_name` to a different value than `db_name` cannot successfully install additional application server instances because SWPM fails to create the corresponding database service in Oracle Clusterware when it calls SAP Host Agent on one of the VM cluster nodes for this task.

To work around this problem, modify the `/oracle/<SID>/sapprof/serviceRac.sh` script on each VM cluster node before running SWPM to install additional application server instances.

For example, make the following change:

## Before:

```
export ORACLE_HOME=/oracle/MFG/19
export ORACLE_SID=MFG
$ORACLE_HOME/bin/srvctl add service -db $ORACLE_SID -service $1 -preferred $2 -available $3 -
tafpolicy BASIC -policy AUTOMATIC -notification TRUE -failovertype SELECT -failovermethod BASIC
-failoverretry 3 -failoverdelay 5
$ORACLE_HOME/bin/srvctl start service -db $ORACLE_SID -service $1
```

## After:

```
export ORACLE_HOME=/oracle/MFG/19
export ORACLE_SID=MFG
export DBUNIQUE_NAME=MFG_LON
$ORACLE_HOME/bin/srvctl add service -db $DBUNIQUE_NAME -service $1 -preferred $2 -available $3 -
tafpolicy BASIC -policy AUTOMATIC -notification TRUE -failovertype SELECT -failovermethod BASIC
-failoverretry 3 -failoverdelay 5
$ORACLE_HOME/bin/srvctl start service -db $DBUNIQUE_NAME -service $1
```

## References

### SAP Notes

- [2614028 - SAP NetWeaver on Oracle Database Exadata Cloud Service](#)
- [2956661 - SAP NetWeaver with Oracle Exadata Cloud@Customer](#)

### Oracle Technical Briefs

- [SAP NetWeaver Application Server ABAP/Java with Oracle Exadata Database Service Based on Oracle Exadata Cloud Infrastructure X11M](#)
- [SAP NetWeaver Application Server ABAP/Java with Oracle Exadata Cloud@Customer X11M](#)

## Connect with us

Call +1.800.ORACLE1 or visit [oracle.com](https://www.oracle.com). Outside North America, find your local office at [oracle.com/contact](https://www.oracle.com/contact).

 [blogs.oracle.com](https://blogs.oracle.com)

 [facebook.com/oracle](https://facebook.com/oracle)

 [twitter.com/oracle](https://twitter.com/oracle)

Copyright © 2025, 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.

Oracle 19c Data Guard and SAP NetWeaver® Application Server ABAP/Java: A Complementary Guide for Oracle Exadata Database Service  
September 2025

Author: Markus Breunig

Contributing Authors: Timm Seitz, Torsten Grambs, Jan Klokkers