

Hybrid Data Guard to Exadata Cloud Services

Production Database on Premises and Disaster Recovery with
Oracle Cloud Infrastructure Exadata Cloud Services

ORACLE WHITE PAPER | JULY 2019





Contents

Introduction	1
Disaster Recovery to the Cloud with Data Guard and Active Data Guard	1
Benefits of Hybrid Standby in the Cloud	2
Enabling Disaster Recovery in Exadata Cloud Services	3
Service Level Requirements	3
Security Requirements	4
Database, OS Environment and Network Prerequisites	5
Using Standby Database to reduce downtime during Planned Maintenance	5
OCI Network Prerequisite	6
Secure Connectivity	6
Public Internet Connectivity	6
On-Premises Network Configuration	7
On-Premises Prerequisites	7
Implement MAA Best Practice Parameter Settings on the Primary Database	7
Validating Connectivity between On-Premises and OCI Exadata Host	8
Deployment Process	8
Verify necessary rpms	8
Step 1: Create Database Using the create database API	8
Step 2: Manually Delete the Database Created By the API	9
Step 3: Copy the Password File to the ExaCS hosts.	10
Step 4: Copying the wallet file to ExaCS.	11
Step 5: Configure static listeners	11



Step 6: Oracle Net Encryption and TNS Entries for Redo Transport	12
Step 7: Create Online and Standby Redo Logs	14
Step 8: Instantiate the Standby Database	15
Step 9: Add RAC instances to tnsnames.ora (RAC only)	16
Step 10: Configure Data Guard broker.	17
Step 11: Remove the static listener.	18
Configuring Client Failover	18
Health Check and Monitoring	18
Oracle MAA Scorecard	18
Monitoring	19
Validate DR Readiness	19
Converting Standby Database to a Snapshot Standby	19
Failover/Switchover to the Cloud	20
Switch back to On-Premises	21
Conclusion	21
Appendix A: Example Output of RMAN Duplicate:	22
Appendix B: MAA Best Practice Parameter Settings	27
Appendix C: Example output for RMAN RESTORE..FROM SERVICE	28
Appendix D: Example create database JSON file	31



Introduction

Oracle's Maximum Availability Architecture (Oracle MAA) is the best practices blueprint for data protection and availability for Oracle databases deployed on private, public or hybrid clouds. Data Guard and Active Data Guard provide disaster recovery (DR) for databases with recovery time objectives (RTO) that cannot be met by restoring from backup. Customers use these solutions to deploy one or more synchronized replicas (standby databases) of a production database (the primary database) in physically separate locations to provide high availability, comprehensive data protection, and disaster recovery for mission-critical data.

An effective disaster recovery plan can be costly due to the need to establish, equip and manage a remote data center. The Oracle Cloud offers a great alternative for hosting standby databases for customers who do not have a DR site or who prefer not to deal with the cost or complexity of managing a remote data center. Existing production databases remain on-premises and standby databases used for DR are deployed on the Oracle Cloud. This mode of deployment is commonly referred to as a hybrid cloud implementation.


Customers may choose to deploy either a Data Guard or an Active Data Guard standby on the cloud depending upon their requirements. While there are some unique considerations to a hybrid cloud DR configuration, it follows the same Oracle MAA best practices as with any Data Guard deployment. This Oracle MAA blueprint details Oracle MAA Best Practices and provides a procedural overview for deploying DR on the Oracle Cloud using Database as a Service. This paper is intended for a technical audience having knowledge of Oracle Database, Data Guard or Active Data Guard, and Oracle Database backup and recovery. This paper also assumes a basic understanding of services offered by the Oracle Cloud Infrastructure.

Disaster Recovery to the Cloud with Data Guard and Active Data Guard

The Oracle Cloud offers an extensive set of cloud services tailored to specific customer requirements: Platform as a Service (PaaS), Infrastructure as a Service (IaaS), and Software as a Service (SaaS). Disaster Recovery (DR) for on-premises systems is deployed using the Exadata Cloud Service (ExaCS).

Data Guard is included in Oracle Database Enterprise Edition (no separate license is required for on-premises systems) and is supported by all editions of OCI DBaaS (Enterprise, High Performance, and Extreme Performance). Please bear in mind standard edition does not support Data Guard. Please search "Licensing Information" in the [Oracle Database Documentation](#) for the selected version to understand the supportability matrix on what is supported in Enterprise Edition versus Standard Edition.

Active Data Guard extends Data Guard capabilities by providing advanced features for data protection and availability as well as offloading read-only workload and fast incremental backups from a production database.



Active Data Guard is included in the Extreme Performance Edition and Exadata Service. When used in a hybrid configuration, Active Data Guard must also be licensed for the on-premises system.

MAA recommends

1. Create an ExaCS standby database target that is symmetric or similar to on-premise primary database to ensure you meet the same performance SLAs after a role transition. Use RAC for RAC, Exadata for Exadata, etc.
2. Ensure network bandwidth is enough to handle peak redo rates
3. Ensure network reliability and security between on-premise to and from OCI cloud
4. Use Active Data Guard for additional auto-block repair, data protection and offloading benefits
5. Use Oracle Transparent Data Encryption (TDE) for both primary and standby databases

Benefits of Hybrid Standby in the Cloud

1. Cloud data center and infrastructure is managed by Oracle
2. Cloud provides basic system life cycle operations including scaling CPU resources
3. Oracle Data Guard provides disaster recovery, data protection and ability to offload activity for higher utilization and return on investment.
4. When configured with MAA practices, RTO of seconds with automatic failover when configured with Data Guard Fast-Start failover and RPO less than a second for Data Guard with ASYNC transport or RPO zero for Data Guard in SYNC or FAR SYNC configurations

Data Guard Life Cycle Management e.g. switchover, failover and reinstate is a manual process in Hybrid Data Guard configurations.

Enabling Disaster Recovery in Exadata Cloud Services

Enabling DR on the cloud requires instantiation of a Data Guard standby database in ExaCS. Once instantiated, Data Guard maintains synchronization between the primary database on-premises and the standby database in the cloud.

- » Utilization of the same Oracle MAA best practices as on-premises deployment. Use of additional Oracle MAA best practices specific to hybrid cloud deployments that are specified in this paper.
- » Ability to switchover (planned events) or failover (unplanned event) production to the standby database in the cloud during planned maintenance or unplanned outages. Once the failed on-premises database is repaired, the ability to automatically resynchronize it with the new production database in the cloud and then switch production back to the on-premises database.

Service Level Requirements

Hybrid cloud deployments are by definition user-managed environments. The administrator must determine service level expectations for availability, data protection, and performance that are practical for a given configuration and application. Service Levels must be established for each of three dimensions relevant to disaster recovery that are applicable to any Data Guard configuration:

- » **Recovery Time Objective (RTO)** describes the maximum acceptable downtime should an outage occur. This includes the time required to detect the outage and to failover both the database and application connections so that service is resumed.
- » **Recovery Point Objective (RPO)** describes the maximum amount of data loss that can be tolerated. Achieving the desired RPO depends upon:
 - » Available bandwidth relative to network volume.
 - » The ability of the network to provide reliable, uninterrupted transmission.
 - » The Data Guard transport method used: either asynchronous for near-zero data loss protection or synchronous for zero data loss protection.
- » **Data Protection:** With Active Data Guard and MAA configuration, customers can configure the most comprehensive [block corruption detection, prevention and auto-repair](#).
- » **Performance:** Database response time may be different after failover if less capacity – compute, memory, I/O, etc, are provisioned at the standby system than in the on-premises production system. This occurs when administrators purposefully under-configure standby resources to reduce cost; accepting reduced service level while in DR mode. MAA best practices recommend configuring symmetrical capacity at both primary and standby so there is no change in response time after failover. Rapid provisioning available with the cloud can enable a middle ground where there is less capacity deployed steady-state, but the new primary is rapidly scaled-up should a failover be required.

Security Requirements

Oracle MAA best practice recommends using Oracle Transparent Data Encryption (TDE) to encrypt both primary and standby databases to ensure all data is encrypted at-rest. Data can be converted during the migration process but it's highly recommended to convert to TDE prior to migration to provide the most secure Data Guard environment. Refer to Oracle Database Tablespace Encryption Behavior in Oracle Cloud (Doc ID 2359020.1) for more information. VPN connection or Oracle Net encryption is also required for encryption-in-flight for any other database payload (e.g. data file or redo headers) that are not encrypted by TDE.

Using TDE to protect data is an important part of improving the security of the system. Users should, however, be aware of certain considerations when using any encryption solution, including:

- » Additional CPU overhead: Encryption requires additional CPU cycles to calculate encrypted and decrypted values. TDE, however, is optimized to minimize the overhead by taking advantage of database caching capabilities and leveraging hardware acceleration within Exadata. Most TDE users see little performance impact on their production systems after enabling TDE. If performance overhead is a concern, please refer to the Oracle Database Advanced Security Guide.
- » Lower data compression: Encrypted data compresses poorly because it must reveal no information about the original plaintext data. Thus, any compression applied to TDE encrypted data will have low compression ratios. Hence, when TDE encryption is used, it is not recommended to use with redo transport. However, when TDE is used in conjunction with Oracle databases compression technologies such as Advanced Compression or Hybrid Columnar Compression, compression is performed before the encryption occurs, and the benefits of compression and encryption are both achieved.
- » Key management: Encryption is only as strong as the key used to encrypt. Furthermore, the loss of the encryption key is tantamount to losing all data protected by that key. If encryption is enabled on a few databases, keeping track of the key and its lifecycle is relatively easy. As the number of encrypted databases grows, managing keys becomes an increasingly difficult problem. For users with a large number of encrypted databases, it is recommended that Oracle Key Vault be used on-premise to store and manage TDE master keys.

Database, OS Environment and Network Prerequisites

TABLE 1: PREREQUISITES

	On-Premises	Oracle Cloud Infrastructure Exadata Cloud Services
Operating System	Linux, Windows or Solaris X86 (My Oracle Support Note 413484.1 for Data Guard cross-platform compatibility)	Oracle Enterprise Linux (64-bit)
Oracle Database*	<ul style="list-style-type: none">» Oracle Database Enterprise Edition 11.2.0.4 (64-bit)» Oracle Database Enterprise Edition 12.1.0.2 (64-bit)» Oracle Database Enterprise Edition 12.2.0.1 (64-bit)» Oracle Database Enterprise Edition 18c (64-bit)» Oracle Database Enterprise Edition 19c (64-bit)	Extreme performance / BYOL
RAC	RAC or non-RAC	RAC or non-RAC
Multitenant ****	For 12.1 and above, primary database has to be a CDB/PDB database.	Multitenant Database Non-CDB
Physical Vs Virtual	Physical or Virtual	Exadata Virtual
Database Size	Any Size	Any size. For shape limits please consult Exadata OCI documentation
TDE Encryption	Mandatory for both primary and standby	Mandatory for both primary and standby

If On-Premises is not already enabled with TDE, please follow the master note [Master Note for Transparent Data Encryption \(TDE\) \(Doc ID 1228046.1\)](#) to enable TDE and create wallet files.

*** Oracle Database version on primary and standby databases must match during initial instantiation. For database software updates that are standby-first compatible, the primary and standby database Oracle Home software can be different. Refer to Oracle Patch Assurance - Data Guard Standby-First Patch Apply (Doc ID 1265700.1)*

Using Standby Database to reduce downtime during Planned Maintenance

There are several options for utilizing a standby database on the cloud for reducing planned downtime of the primary production database:

Standby-first Patch Apply

Many patches may be applied first to a physical standby for thorough validation. Customers who wish to minimize downtime will frequently patch the standby first, then switch production to the standby database, and then patch the original primary. If the primary and standby are RAC and the software update is RAC rolling, a switchover is not required; however it is still recommended to update the software on the standby-first for additional validation and protection. Data Guard physical replication is supported between primary and standby running at mixed patch versions for patches that are standby-first eligible (documented in the patch readme). The customer may also choose to run for a period of time with mixed patch versions between primary and standby to enable fast fallback to the unpatched version should there be any unanticipated problems with the patch. See My Oracle Support Note 1265700.1, "Oracle Patch Assurance - Data Guard Standby-First Patch Apply" for more details on patches eligible for the standby-first process.

Database Rolling Upgrade

Another beneficial use case for standby in Exadata OCI is for database rolling upgrade to reduce downtime when upgrading to new database Oracle releases that are not standby-first compatible. The transient logical process is used in Oracle 11g and Oracle 12c to temporarily convert a physical standby database to a logical standby, upgrade the logical standby to the new version, validate and when ready execute a Data Guard switchover. After the switchover completes, the original primary database is converted to a synchronized physical standby also operating at the new release. Refer to Oracle 11g [Database Rolling Upgrades Made Easy](#) or Oracle 12c [DBMS_Rolling](#) for more information. A more efficient database rolling upgrade process using the standby database exists for Data Guard environments 12.2 and higher. Refer to [Using DBMS_ROLLING to Perform a Rolling Upgrade](#) section in the Data Guard documentation.

OCI Network Prerequisite

Data transfers from on-premises to Oracle Cloud use the public network, VPN and/or the high bandwidth option provided by Oracle FastConnect.

In a Data Guard configuration, the primary and standby must be able to communicate bi-directionally. This requires additional network configuration to allow access to ports between the systems.

Secure Connectivity

There are two options to privately connect your cloud network to the on-premises network, IPSec VPN and FastConnect. Each of these methods require a Dynamic Routing Gateway (DRG) to connect to your Virtual Cloud Network (VCN). Details for creating a DRG can be found in the [documentation at this link](#).

IPSec VPN

IPSec stands for Internet Protocol Security or IP Security. IPSec is a protocol suite that encrypts the entire IP traffic before the packets are transferred from the source to the destination. For an overview of IPSec in OCI review the documentation [here](#).

FastConnect

Oracle Cloud Infrastructure FastConnect provides a method to create a dedicated, private connection between your data center and Oracle Cloud Infrastructure. FastConnect provides higher-bandwidth options and a more reliable and consistent networking experience compared to internet-based connections. More details on FastConnect can be reviewed [here](#).

Public Internet Connectivity

Connectivity between OCI and on-premises can be achieved through the public internet as well. This method is not by default secure and additional steps must be taken to secure transmissions. The remainder of this whitepaper assumes public internet connectivity.

By default, cloud security rule for port 1521 is disabled. Also, this default pre-configured port in the cloud VM/BM has open access from the public internet.

1. If Virtual Cloud Network for the standby database doesn't have Internet Gateway, you have to add an Internet Gateway. The below document provides how to create internet gateway:
<https://docs.us-phoenix-1.oraclecloud.com/Content/Network/Tasks/managingIGs.htm>
2. Edit Ingress and Egress rule in security list to connect from/to the on-premise database. The link below provides additional information.
<https://docs.us-phoenix-1.oraclecloud.com/Content/Network/Concepts/securitylists.htm>

On-Premises Network Configuration

In a Data Guard configuration, information is transmitted in both directions between primary and standby databases. This requires basic configuration, network tuning and opening of ports at both primary and standby databases.

On-Premises Prerequisites

The following prerequisites must be met before instantiating the standby database:

- » Name resolution to the OCI Exadata hosts needs to be configured. This can be done either through a static file like /etc/hosts, or configuring the on-premises DNS to properly resolve the public IP address of the OCI instance. Also, the on-premises Firewall will need to have Access Control Lists properly configured to allow SSH and Oracle Net to be accessed from the on-premises system to OCI.
- » Since Data Guard in a DR situation requires access from the cloud instance to the on-premises database, the primary database listener port must be opened with restricted access from the cloud IP addresses using features like iptables etc. Since every corporation has different network security policies, the network administrator will need to perform operations similar to the cloud-side network configuration shown in preceding sections
- » Prompt-less SSH from the OCI Exadata to the On-Premises machine. This is configured both for on-premises to the Cloud during the provisioning process and from the Cloud to on-premises.
- » The configuration of the On-Premises firewall to allow inbound ssh connectivity from the Exadata OCI to the On-Prem machine.
- » The Oracle Home for the on-premises database must be the same Oracle patchset as the standby database. If the OCI environment is on a different bundle patch level or PSU compared to the on-premise, it is recommended to patch the source environment to the same database bundle patch level as the database home in the cloud environment. (The command "\$ORACLE_HOME/OPatch/patch lspatches" can be executed to check the one-off patches installed on both Source and Target environments)
- » The steps outlined in this document assume that the on-premises primary database is not already part of an existing Data Guard broker configuration. If there is an existing broker configuration for the on-premises database it is assumed that the administrator has prior knowledge of the broker and knows how to add the new standby database to an existing broker configuration. A value other than 'NOCONFIG' for the following query implies an existing broker configuration.

```
SQL> select decode(count(1),0,'NOCONFIG') from v$DG_BROKER_CONFIG;
```


- » Use the default listener named LISTENER. The steps outlined in this document assume the default listener name LISTENER is used. To verify run the following command from the on-premises machine. The expected result is shown.

```
$lsnrctl show current_listener | grep 'Current Listener'
Current Listener is LISTENER
```

- » Verify the listener port by running the following command from the on-premises machine. The expected result is shown.

```
$lsnrctlstat | grep 'Connecting to'
Connecting to (ADDRESS=(PROTOCOL=tcp)(HOST=)(PORT=(1521)))
```

Implement MAA Best Practice Parameter Settings on the Primary Database



See Appendix B for a list of best practices. Completing this process on the primary database before instantiation is recommended. Especially configuring the redo logs which will be duplicated during the process.

Validating Connectivity between On-Premises and OCI Exadata Host

Once all the steps above are implemented successfully, run the command below to validate if the connection looks good from all sources to all targets and vice versa. If telnet is show success, proceed to the next step.

ON ON-PREMISE HOST

```
[root@onpremise1 ~]# telnet xxx.xxx.xxx.xxx 1521
Trying xxx.xxx.xxx.xxx...
Connected to xxx.xxx.xxx.xxx.
Escape character is '^]'.
^C^]q
telnet> q
Connection closed.
```

ON EXACS HOSTS

```
[root@oci2 ~]# telnet xxx.xxx.xxx.xxx 1521
Trying xxx.xxx.xxx.xxx...
Connected to xxx.xxx.xxx.xxx.
Escape character is '^]'.
^]q
telnet> q
Connection closed.
```

Deployment Process

This procedure is similar to migrating the database from on-premise to OCI.

<https://docs.us-phoenix-1.oraclecloud.com/Content/Database/Tasks/mig-rman-duplicate-active-database.htm>

PREREQUISITES

It is recommended to create the database with the same name as the on-premise database. While the db_name should be the same, the db_unique_name must be different from the primary system. The patch level of the source database's Oracle home should be the same as the database on Exa-OCI.

Verify necessary rpms

This process requires a minimum dbaastools rpm version of 18.2.3.2.0_190618 but updating to the latest dbaastools rpms are always recommended. Confirm the current installed version with the following command.

```
$ rpm -qa | grep dbaastools
dbaastools_exa-1.0-1+18.2.3.2.0_190618.1843.x86_64
```

Apply the latest tools rpm with the following:

```
(as root) # dbaascli patch tools apply LATEST
```

Step 1: Create Database Using the create database API

Using the [create database API](#), create a new database. Be sure to:

- set `dbname` to the `db_name` parameter of the primary database
- include the `db_unique_name` parameter with a unique name different than the `dbname`
- set `version` to the version of the primary database from the list of available versions in the link.

Step 2: Manually Delete the Database Created By the API

Once the database is provisioned and ready, perform the below operations to delete the provisioned database files and restore the on-premises database using RMAN. The steps to initiate the RMAN duplicate will be described later in this document.

To delete the provisioned database use the manual method of removing the database files from ASM diskgroups. Do not use DBCA as this will also remove the `srvctl` registration as well as the `oratab` entries which should be retained for the standby.

To manually delete the database on the Exadata Cloud Services hosts, run the steps below

1. Using the standby database `db_unique_name` defined in the create database json file.

2. Create a script to remove all database files

```
SQL> set heading off linesize 999 pagesize 0 feedback off trimspool on
SQL> spool /tmp/files.lst
SQL> select 'asmcmd rm ' || name from v$datafile
union all
select 'asmcmd rm ' || name from v$tempfile
union all
select 'asmcmd rm ' || member from v$logfile;
SQL> spool off

SQL> create pfile='/tmp/<standby DB_UNIQUE_NAME>.pfile' from spfile; #Backup spfile
$chmod 777 /tmp/files.lst
```

3. Shutdown the database

First collect the configuration of the database for future reference:

```
$ srvctl config database -d <db_unique_name> > /tmp/<standby db_unique_name>.config
```

Then stop the database:

```
$ srvctl stop database -d <db_unique_name> -o immediate
```

4. Remove database files

Remove the existing data files, log files and tempfile. The password file will be replaced and the spfile will be reused.

As grid user (sudo from `opc` user to `grid` user)

Edit /tmp/files.lst created previously to remove any unneeded lines from sqlplus, leaving only lines beginning with 'asmcmd'.

Save and execute the script

```
[grid@<host> ~]$ . /tmp/files.lst
```

All files for the starter database have now been removed.

Step 3: Copy the Password File to the ExaCS hosts.

Copy on-premises password file to ExaCS hosts.

Check password file location

If Oracle Clusterware is running on the on-Premises host, check the password file location with "srvctl config database -db <db_name>". If Oracle Clusterware isn't configured or password file location is null, password file exists in default location(\$ORACLE_HOME/dbs/orapwd<ORACLE_SID>)

```
$ srvctl config database -db testdbname
Database unique name: testdbname
Database name:
Oracle home: /u02/app/oracle/product/18.0.0/dbhome_2
Oracle user: oracle
Spfile: +DATA/testdbname/spfiledbtestdbname.ora
Password file: +DATA/testdbname/PASSWORD/orapwd<sid> <===== password file location
Domain: domainname.xxxx.xxxx
```

Copy password file on-premises to ExaCS hosts

Copy password file to all ExaCS nodes.

If password file location is non-ASM, copy the file as below.

```
$ scp -i <ssh key> $ORACLE_HOME/dbs/orapw<SID> opc@<Public-IP-OCI-HOST>:/tmp
```


If password file location is ASM, switch user to "grid" or the ASM owner, source the environment variables and then copy the password file as below.

```
$ sudo su - grid
$ export ORACLE_SID=<ASM ORACLE_SID>
$ export ORACLE_HOME=<GRID_HOME>
$ asmcmd
ASMCMD> cd +<DISKGROUP_NAME>/<DB_UNIQUE_NAME>/PASSWORD
ASMCMD> cp orapw<SID> /tmp
copying +DATA/testdbname/PASSWORD/orapw<sid> -> /tmp/orapw<sid>

scp -i <ssh key> /tmp/orapw<SID> opc@<Public-IP-OCI-HOST>:/tmp
```

Place the password file in the proper ExaCS location for the standby database.

As opc user on ExaCS host



```
$ chmod 777 /tmp/<password file name>
$ sudo su - grid
```

Place the password file in ASM:

Note use `srvctl config database -db <standby DB_UNIQUE_NAME>` to find current passwordfile location

As grid user:

```
$ asmcmd pwcop --dbunique <standby DB_UNIQUE_NAME> /tmp/pwdvictor.858.1001423135
+DATA1/victorsb/PASSWORD/passwd -f
```

NOTE: Error ASMCMD-9453: failed to register password file as a CRS resource can be ignored if you are reusing the previously registered password file location. :

Step 4: Copying the wallet file to ExaCS.

Retrieve the TDE wallet location from the on-premises database with the following query:

```
SQL> select WRL_PARAMETER from v$encryption_wallet;
```

WRL_PARAMETER

```
-----
+DATA1/<DB_NAME>/TDE/
```

Copy the ewallet.p12 and cwallet.sso files from on-Premises to the above directory on ExaCS node1. Note: it may be necessary to copy the files out of ASM to /tmp as with the password file.

ON ON-PREMISE HOST


```
scp -i ~/<ssh_key> <PATH>/ewallet.p12 opc@<Public-IP-OCI-HOST1>:/tmp
scp -i ~/<ssh_key> <PATH>/cwallet.sso opc@<Public-IP-OCI-HOST1>:/tmp
```

Remove old wallet files in `/var/opt/oracle/dbaas_acfs/<standby db_name>/tde_wallet`. Execute the commands only on node1. (ExaCS nodes share storage under `/var/opt/oracle/dbaas_acfs` using ACFS)

ON EXACS HOST NODE1

```
$ chmod 777 /tmp/ewallet.p12
$ chmod 777 /tmp/cwallet.sso
$ sudo su - oracle
$ cp /tmp/ewallet.p12 /var/opt/oracle/dbaas_acfs/<standby db_name>/tde_wallet/
$ cp /tmp/cwallet.sso /var/opt/oracle/dbaas_acfs/<standby db_name>/tde_wallet/
$ chmod 600 /var/opt/oracle/dbaas_acfs/<standby db_name>/tde_wallet/*wallet*
```

Step 5: Configure static listeners



A static listener is needed for initial instantiation of a standby database. The static listener enables remote connections to an instance while the database is down in order to start a given instance. See MOS 1387859.1 for additional details.

As the grid user or software owner, add the following entry to the listener.ora on the ExaCS node 1 after replacing the variables. The listener.ora is located in <grid home>/network/admin

```
LISTENER.ORA
SID_LIST_LISTENER =
  (SID_LIST = (SID_DESC =
    (GLOBAL_DBNAME = <DB_UNIQUE_NAME of the ExaCS database>) (ORACLE_HOME = <Local
Oracle Home>)
    (SID_NAME = <ORACLE SID of the local instance>)
  ))
```

For 11.2 configurations, a static listener is also required for Data Guard Broker. Add the following entry to the listener.ora on premises after replacing the variables. For RAC this must be done on both nodes.

```
SID_LIST_LISTENER =
  (SID_LIST = (SID_DESC =
    (GLOBAL_DBNAME = <DB_UNIQUE_NAME of the ExaCS database>_DGMGRL)
    (ORACLE_HOME = <Local Oracle Home>) (SID_NAME = <ORACLE SID of the local
instance>)
  ))
```

Finally reload the listener (as grid user):

```
$ORACLE_HOME/bin/lsnrctl reload
```

Start the Standby Instance

```
srvctl start instance -d <standby DB_UNIQUE_NAME> -i <standby instance name> -o mount
```

Step 6: Oracle Net Encryption and TNS Entries for Redo Transport

To protect from plaintext or redo from unencrypted tablespaces from being visible on the WAN place the following entries in the sqlnet.ora file on all on-premises and cloud machines which are located in \$ORACLE_HOME/network/admin.

```
SQLNET.ORA ON ON-PREMISE HOST(S)
SQLNET.ENCRYPTION_SERVER=REQUIRED
SQLNET.CRYPTO_CHECKSUM_SERVER=REQUIRED
SQLNET.ENCRYPTION_TYPES_SERVER=(AES256,AES192,AES128)
SQLNET.CRYPTO_CHECKSUM_TYPES_SERVER=( SHA1 )
SQLNET.ENCRYPTION_CLIENT=REQUIRED
SQLNET.CRYPTO_CHECKSUM_CLIENT=REQUIRED
SQLNET.ENCRYPTION_TYPES_CLIENT=(AES256,AES192,AES128)
SQLNET.CRYPTO_CHECKSUM_TYPES_CLIENT=( SHA1 )
```

Entries for each database are needed in both primary and standby `tnsnames.ora` files for proper redo transport. Use the following example, replacing **bolded** values with values relevant to the configuration.

NOTE: The primary database may already have a TNS entry in the on-premises `tnsnames.ora` with a server name for the HOST. In this case, simply change the server name in that entry to use the IP address for the host instead.

NOTE: IP addresses are used since there is no DNS between on-premises and cloud environments to resolve server names to IP addresses.

NOTE: RAC configurations cannot resolve the scan listener name thus an address list must be used listing all nodes. Those will be added later, the initial instanton should only list one IP in the tns entries to ensure RMAN is always connecting to the same nodes.

TNSNAMES.ORA ON ON-PREMISE HOST

```
<standby db_unique_name> =
(DESCRIPTION =
  (SDU=65536) (RECV_BUF_SIZE=134217728)
  (SEND_BUF_SIZE=134217728)
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP)(HOST = <standby IP address>)(PORT = {1521|<port#>}))
  )
  (CONNECT_DATA =
    (SERVER = DEDICATED)
    (SERVICE_NAME = <DB_UNIQUE_NAME of the ExaCS database>)
    (UR=A)
  ))
```


TNSNAMES.ORA ON EXACS HOSTS

```
<primary db_unique_name> =
(DESCRIPTION =
  (SDU=65536) (RECV_BUF_SIZE=134217728)
  (SEND_BUF_SIZE=134217728)
  (ADDRESS_LIST =
    (ADDRESS = (PROTOCOL = TCP)(HOST = <primary IP address>)(PORT = {1521|<port#>}))
  )
  (CONNECT_DATA =
    (SERVER = DEDICATED)
    (SERVICE_NAME = <primary db_unique_name>[.<primary domain name>])
    (UR=A)
  ))
```

Set TCP socket size

Check the TCP socket sizes for the on premises system as well as the cloud instance with the following command run as root.

```
ON ON-PREMISE HOST(S)
# /sbin/sysctl -a | egrep net.core.[w,r]mem_max
```

```
net.core.wmem_max = 4194304
net.core.rmem_max = 4194304
```

ON EXACS HOSTS

```
# /sbin/sysctl -a | egrep net.core.[w,r]mem_max
net.core.wmem_max = 4194304
net.core.rmem_max = 4194304
```

If necessary, adjust all socket size maximums to 128MB or 134217728. For on premises systems consult your operating system guide for details about how to accomplish this. For the cloud instance edit the `/etc/sysctl.conf` file settings for `net.core.wmem_max` and `net.core.rmem_max`. If the values between on premises and OPC do not match, the network protocol will negotiate the lower of the two values. Therefore, the values between sites is not required to match though that is recommended in order to attain optimal transport performance.

```
net.core.rmem_max = 134217728 net.core.wmem_max = 134217728
```

Step 7: Create Online and Standby Redo Logs

On the on-premises primary database, add online redo logs to each thread so there are at least three groups per thread. All log files group should be of the same size.

```
SQL> select thread#,group#,bytes from v$log;
```

THREAD#	GROUP#	BYTES
1	1	4294967296
1	2	4294967296
2	3	4294967296
2	4	4294967296

```
SQL> alter database add logfile thread 1 group 5 size 4G;
```

Database altered.


```
SQL> alter database add logfile thread 2 group 6 size 4G;
```

Database altered.

```
SQL> select thread#,group#,bytes from v$log;
```

THREAD#	GROUP#	BYTES
1	1	4294967296
1	2	4294967296
2	3	4294967296
2	4	4294967296
1	5	4294967296
2	6	4294967296

Now add an equal number of standby redo logs per thread as there are online redo logs. Standby Redo Logs should NOT be multiplexed and should be on the DATA disk group to ensure best performance.



```
SQL> alter database add standby logfile thread 1 group 7 ('+DATA1') size 4G, group 8
('+DATA1') size 4G, group 9 ('+DATA1') size 4G;
```

Database altered.

```
SQL> alter database add standby logfile thread 2 group 10 ('+DATA1') size 4G, group 11
('+DATA1') size 4G, group 12 ('+DATA1') size 4G;
```

Database altered.

```
SQL> select thread#,group#,bytes from v$standby_log;
```

THREAD#	GROUP#	BYTES
1	7	4294967296
1	8	4294967296
1	9	4294967296
2	10	4294967296
2	11	4294967296
2	12	4294967296

Step 8: Instantiate the Standby Database

The standby database can be created from the active primary database or from a backup of the primary database. This section describes the method duplicating from the active primary database using Oracle 12.1 or higher feature RMAN 'RESTORE...FROM SERVICE'.

Backups can also be used to instantiate databases and may be more efficient depending on the size of the database and transfer rate between the systems.

RDBMS 11.2 does not support RMAN RESTORE...FROM SERVICE. Backups based duplication or RMAN DUPLICATE must be used

For details see the [documentation](#).


For additional details regarding standby instantiation see [MOS 2275154.1 Creating a Physical Standby Database in an 11.2, 12.1, 12.2 or later environment](#)

Start the standby instance (one instance for RAC)

```
$ srvctl stop database -d <standby DB_UNIQUE_NAME> -o immediate
$ rman target /
RMAN> startup nomount
```

```
RMAN> restore standby controlfile from service 'primary';
```

```
RMAN> alter database mount;
```



```
RMAN> CONFIGURE DEVICE TYPE DISK PARALLELISM 4; ← sets the degree or parallelism (number
of channels) and can be altered depending on bandwidth.
```

```
RMAN> restore database from service 'primary' section size 5G;
```

```
RMAN> shutdown immediate
```

Restart the standby database

```
$ srvctl start database -d <standby DB_UNIQUE_NAME> -o mount
```

Clear all online and standby redo logs

```
$ sqlplus "/ as sysdba"
SQL> set pagesize 0 feedback off linesize 120 trimspool on
SQL> spool /tmp/clearlogs.sql
SQL> select distinct 'alter database clear logfile group '||group#||';' from
v$logfile;
SQL> spool off
SQL> @/tmp/clearlogs.sql

SQL> select member from v$logfile;
```

NOTE: All redo logs should be on the DATA disk group in the standby DB_UNIQUE_NAME directory

```
SQL> alter database flashback on;
```

Step 9: Add RAC instances to tnsnames.ora (RAC only)

Since the scan name cannot be resolved in either direction in a hybrid configuration, and address list must be configured to provide high availability.

NOTE: List the addresses in the address list in a different order on each node to balance redo transport across nodes

TNSNAMES.ORA ON ON-PREMISE HOST

```
<standby db_unique_name> =
  (DESCRIPTION=
    (SDU=65536) (RECV_BUF_SIZE=134217728)
    (SEND_BUF_SIZE=134217728)
    (ADDRESS_LIST=
      (FAILOVER=on)
      (CONNECT_TIMEOUT=3) (RETRY_COUNT=3)
      (ADDRESS = (PROTOCOL = TCP)(HOST = <standby node1 IP address>)(PORT = {1521|<port#>}))
```

```

)
  (ADDRESS = (PROTOCOL = TCP)(HOST = <standby node2 IP address>)(PORT = {1521|<port#>}))
)
)
(CONNECT_DATA=
  (SERVER=DEDICATED)
  (SERVICE_NAME= standby db_unique_name>[.<standby domain name>])
)
)

TNSNAMES.ORA ON EXACS HOSTS
<primary db_unique_name> =
  (DESCRIPTION=
    (SDU=65536) (RECV_BUF_SIZE=134217728)
    (SEND_BUF_SIZE=134217728)
    (ADDRESS_LIST=
      (FAILOVER=on)
      (CONNECT_TIMEOUT=3)(RETRY_COUNT=3)
      (ADDRESS = (PROTOCOL = TCP)(HOST = <primary node1 IP address>)(PORT =
{1521|<port#>}))
    )
      (ADDRESS = (PROTOCOL = TCP)(HOST = <primary node2 IP address>)(PORT =
{1521|<port#>}))
    )
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = <primary db_unique_name>[.<primary domain name>])
    )
  ))

```

Step 10: Configure Data Guard broker.

Enable the dg_broker_config_file parameter on primary and standby database.

NOTE: For ASM, place broker config files on separate disk groups. For RAC, broker config files must be on shared storage.

ON ON-PREMISE AND EXACS HOSTS

```

SQL> alter system set dg_broker_config_file1='+DATAC1/<db_unique_name>/dr1.dat';
SQL> alter system set dg_broker_config_file2='+RECO1/<db_unique_name>/dr2.dat';

```

Start the Data Guard Broker Process on primary and standby database


ON ON-PREMISE AND EXACS HOSTS

```

SQL> alter system set dg_broker_start=true;
SQL> show parameter dg_broker_start

```

NAME	TYPE	VALUE
------	------	-------



```
dg_broker_start          boolean      TRUE
```

```
SQL> select pname from gv$process where pname like 'DMON%';
```

```
PNAME
-----
DMON
DMON
```

Register the database via DGMGRL on primary site

```
ON ON-PREMISE HOST
$ dgmgrl sys/<sys password>@<net service name for primary database>
DGMGRL> CREATE CONFIGURATION <configuration_name> AS PRIMARY DATABASE IS <primary
database_name> CONNECT IDENTIFIER IS <Net Service name for primary database>;
DGMGRL> ADD DATABASE <standby database_name> AS CONNECT IDENTIFIER IS <Net Service
name for standby database> MAINTAINED AS PHYSICAL;
DGMGRL> enable configuration;
```

Step 11: Remove the static listener.

As grid user on the ExaCS node 1, remove the static listener entries created in step 5 and reload the listener.

Configuring Client Failover

Automating client failover, the process by which clients are reconnected to the active primary database after a failure, includes relocating database services to the new primary database as part of a Data Guard failover, notifying clients that a failure has occurred in order to break them out of TCP timeout, and redirecting clients to the new primary database. Configuration details are thoroughly covered in the papers MAA Best Practices for Client Failover for Oracle [Database 11g](#) and for Oracle [Database 12c](#). Customers using Oracle Database 12c and higher may also choose to configure [application continuity](#) for the most comprehensive continuous availability of applications. Please consult these papers and configure your environment appropriately.

Health Check and Monitoring

After the standby is instantiated, a health check should be performed to ensure the Data Guard databases (primary and standby) are compliant with Oracle MAA best practices. It is also advisable to perform the health check on a monthly basis as well as before and after database maintenance. There are several methods for checking the health of a Data Guard configuration:

Oracle MAA Scorecard

Oracle provides several automated health check tools that can be downloaded from My Oracle Support specific for the type of hardware platform:



» [exachk](#) applicable to Exadata Database Machine (suitable for Exadata Cloud Service)

Each of the automated checks include an Oracle MAA Scorecard that reports on a number of key Data Guard configuration best practices in addition to many other checks.

Oracle strongly recommends the use of these automated tools for comprehensive health check of not only the Data Guard configuration but the system as a whole. The health checks are regularly updated with current information. Be sure to download the latest version of the health checks applicable to your platform.

Monitoring

Regular monitoring of the Data Guard configuration is not provided in a Hybrid Data Guard Configuration and must be done manually. Refer to MOS note [Monitoring a Data Guard Configuration \(Doc ID 2064281.1\)](#) for MAA best practice recommendations for monitoring.

Validate DR Readiness

Best practice is to use Active Data Guard to offload read-only workload to the standby database to provide continuous, application-level validation that the standby is ready for production. This provides a level of assurance in addition to continuous Oracle block-level validation performed by Data Guard apply processes. It is also a best practice to periodically place the standby in read/write mode (using Data Guard Snapshot Standby) to validate its readiness to support read-write production workloads. A snapshot standby may also be used for a final level of pre-production functional and performance testing of patches and upgrades since the DR system is sized similarly to the production system. A Snapshot Standby continues to receive redo from the primary database where it is archived for later use, thus providing data protection at all times. Recovery time (RTO), however, will be extended by the amount of time required to convert the Snapshot Standby back to the standby database if a failover is required while testing is in progress. Note that additional storage is required for the fast recovery area when a standby is in snapshot mode (to hold archived redo received from the primary production database for later use and current redo and flashback logs generated by the snapshot standby). Steps for converting a standby to a snapshot standby and back are listed in the section below. Please refer to Oracle documentation for additional details on Data Guard Snapshot Standby. Optionally, you may perform an actual switchover or failover operation to the cloud for a complete end-to-end DR test; for more details see [Failover/Switchover to the Cloud](#).

Converting Standby Database to a Snapshot Standby

A snapshot standby is a fully updatable standby database that is created from a physical standby database. On snapshot standby databases, redo data is received but not applied until the snapshot standby database is converted back to a physical standby database.

The benefits of using a snapshot standby database include the following:

1. It provides an exact replica of a production database for development and testing purposes while maintaining data protection at all times. You can use the Oracle Real Application Testing option to capture primary database workload and then replay it for test purposes on the snapshot standby.
2. It can be easily refreshed to contain current production data by converting to a physical standby and resynchronizing.

Follow the steps below to convert a physical standby database to a snapshot standby

Convert the standby to a snapshot standby and validate

Via Data Guard broker issue the following commands

```
DGMGRL> convert database 'stby' to snapshot standby;
DGMGRL> SHOW CONFIGURATION;
Configuration - DRSolution
Protection Mode: MaxPerformance Databases:
prmy - Primary database stby - Snapshot standby database
Fast-Start Failover: DISABLED
Configuration Status: SUCCESS
```

NOTE: A snapshot standby database cannot be the target of a switchover or failover. A snapshot standby database must first be converted back into a physical standby database before performing a role transition to it.

Convert the snapshot standby back into a physical standby database

Via Data Guard broker issue the following commands

```
DGMGRL> CONVERT DATABASE 'stby' to PHYSICAL STANDBY;
```

Failover/Switchover to the Cloud


At any time you can manually execute a Data Guard switchover (planned event) or failover (unplanned event). Customers may also choose to automate Data Guard failover by configuring Fast-Start failover. Switchover and failover reverse the roles of the databases in a Data Guard configuration – the standby in the cloud becomes primary and the original on-premises primary becomes a standby database. Refer to Oracle MAA Best Practices for additional information on Data Guard role transitions.

Switchovers are always a planned event that guarantees no data is lost. To execute a switchover perform the following in Data Guard Broker

```
DGMGRL> validate database stby;
Database Role: Physical standby database Primary Database: pri
Ready for Switchover: Yes
Ready for Failover: Yes (Primary Running)
DGMGRL> switchover to <target standby>;
```

A failover is an unplanned event that assumes the primary database is lost. The standby database is converted to a primary database immediately; after all available redo from the primary has been applied. After a failover the old primary database must be reinstated as a physical standby which is made simpler with flashback database and Data Guard broker enabled. To execute a failover and reinstatement execute the following in Data Guard Broker.

```
DGMGRL> failover to stby;
Performing failover NOW, please wait...
Failover succeeded, new primary is "stby"
Execute startup mount on one instance of the old primary before reinstating.
SQL> shutdown abort
SQL> startup mount
DGMGRL> reinstate database pri
Reinstating database "pri", please wait...
```



For more information on role transitions using the Data Guard Broker see the broker documentation for Oracle Database 11g or 12c.

Switch back to On-Premises

The same role transition procedure mentioned in the failover/switchover process is applied again when you are ready to move production back to the on-premises database.

Conclusion

Hybrid Data Guard using Exadata Cloud Services systems is an economical method to achieve Disaster Recovery readiness. Utilizing Maximum Availability Architecture best practices ensures the best solution for data protection and availability.

Appendix A: Example Output of RMAN Duplicate:

```
RMAN> duplicate target database for standby from active database ;
```

```
Starting Duplicate Db at 11-JAN-18
using target database control file instead of recovery catalog
allocated channel: ORA_AUX_DISK_1
channel ORA_AUX_DISK_1: SID=771 device type=DISK
```


```
contents of Memory Script:
{
  backup as copy reuse
    targetfile '+DATA1/racddb/PASSWORD/passwd' auxiliary format
    '/u01/app/oracle/product/18.1.0.0/dbhome_1/dbs/orapwracddb' ;
}
executing Memory Script
```

```
Starting backup at 11-JAN-18
allocated channel: ORA_DISK_1
channel ORA_DISK_1: SID=785 instance=racddb1 device type=DISK
Finished backup at 11-JAN-18
```

```
contents of Memory Script:
{
  sql clone "alter system set  control_files =
    ''+RECO/RACDDB_PHX34Z/CONTROLFILE/current.263.965119865'' comment=
    ''Set by RMAN'' scope=spfile";
  backup as copy current controlfile for standby auxiliary
format '+RECO/RACDDB_PHX34Z/CONTROLFILE/current.264.965119865';
  sql clone "alter system set  control_files =
    ''+RECO/RACDDB_PHX34Z/CONTROLFILE/current.264.965119865'' comment=
    ''Set by RMAN'' scope=spfile";
  shutdown clone immediate;
  startup clone nomount;
}
executing Memory Script
```

```
sql statement: alter system set  control_files
=    ''+RECO/RACDDB_PHX34Z/CONTROLFILE/current.263.965119865'' comment= ''Set by
RMAN'' scope=spfile
```

```
Starting backup at 11-JAN-18
using channel ORA_DISK_1
channel ORA_DISK_1: starting datafile copy
copying standby control file
```



```
output file name=+RECO/RACEDB_PHX34Z/CONTROLFILE/current.264.965119865
```

```
tag=TAG20180111T085105
```

```
channel ORA_DISK_1: datafile copy complete, elapsed time: 00:00:03
```

```
Finished backup at 11-JAN-18
```

```
sql statement: alter system set control_files
```

```
= '+RECO/RACEDB_PHX34Z/CONTROLFILE/current.264.965119865' comment= 'Set by  
RMAN' scope=spfile
```

```
Oracle instance shut down
```

```
connected to auxiliary database (not started)
```

```
Oracle instance started
```

```
Total System Global Area 15032385536 bytes
```

```
Fixed Size 3728304 bytes
```

```
Variable Size 2348813392 bytes
```

```
Database Buffers 12616466432 bytes
```

```
Redo Buffers 63377408 bytes
```

```
contents of Memory Script:
```


```
{  
    sql clone 'alter database mount standby database';  
}
```

```
executing Memory Script
```

```
sql statement: alter database mount standby database
```

```
contents of Memory Script:
```

```
{  
    set newname for clone tempfile 1 to new;  
    set newname for clone tempfile 2 to new;  
    set newname for clone tempfile 4 to new;  
    switch clone tempfile all;  
    set newname for clone datafile 1 to new;  
    set newname for clone datafile 2 to new;  
    set newname for clone datafile 3 to new;  
    set newname for clone datafile 4 to new;  
    set newname for clone datafile 5 to new;  
    set newname for clone datafile 6 to new;  
    set newname for clone datafile 7 to new;  
    set newname for clone datafile 8 to new;  
    set newname for clone datafile 9 to new;  
    set newname for clone datafile 10 to new;  
    backup as copy reuse
```



```
datafile 1 auxiliary format new
datafile 2 auxiliary format new
datafile 3 auxiliary format new
datafile 4 auxiliary format new
datafile 5 auxiliary format new
datafile 6 auxiliary format new
datafile 7 auxiliary format new
datafile 8 auxiliary format new
datafile 9 auxiliary format new
datafile 10 auxiliary format new
;
sql 'alter system archive log current';
}
```

executing Memory Script

executing command: SET NEWNAME

executing command: SET NEWNAME

executing command: SET NEWNAME

renamed tempfile 1 to +DATA in control file

renamed tempfile 2 to +DATA in control file

renamed tempfile 4 to +DATA in control file

executing command: SET NEWNAME

executing command: SET NEWNAME

executing command: SET NEWNAME

executing command: SET NEWNAME

executing command: SET NEWNAME

executing command: SET NEWNAME

executing command: SET NEWNAME

executing command: SET NEWNAME

executing command: SET NEWNAME

executing command: SET NEWNAME

Starting backup at 11-JAN-18

using channel ORA_DISK_1

channel ORA_DISK_1: starting datafile copy

input datafile file number=00009 name=+DATAAC1/racddb/tokyo/tokyo_sysaux.dbf

output file

name=+DATA/RACCDDB_PHX34Z/627B610857F86EA9E0537701000A1DB6/DATAFILE/sysaux.266.9651199

11 tag=TAG20180111T085150

channel ORA_DISK_1: datafile copy complete, elapsed time: 00:00:45

channel ORA_DISK_1: starting datafile copy

input datafile file number=00001 name=+DATAAC1/racddb/DATAFILE/system.dbf


output file name=+DATA/RACCDDB_PHX34Z/DATAFILE/system.265.965119957

tag=TAG20180111T085150

channel ORA_DISK_1: datafile copy complete, elapsed time: 00:00:25



```
channel ORA_DISK_1: starting datafile copy
input datafile file number=00003 name=+DATAC1/racddb/DATAFILE/sysaux.dbf
output file name=+DATA/RACCDDB_PHX34Z/DATAFILE/sysaux.272.965119983
tag=TAG20180111T085150
channel ORA_DISK_1: datafile copy complete, elapsed time: 00:00:25
channel ORA_DISK_1: starting datafile copy
input datafile file number=00005 name=+DATAC1/racddb/DATAFILE/undotbs1.dbf
output file name=+DATA/RACCDDB_PHX34Z/DATAFILE/undotbs1.263.965120007
tag=TAG20180111T085150
channel ORA_DISK_1: datafile copy complete, elapsed time: 00:00:25
channel ORA_DISK_1: starting datafile copy
input datafile file number=00006 name=+DATAC1/racddb/DATAFILE/undotbs2.dbf
output file name=+DATA/RACCDDB_PHX34Z/DATAFILE/undotbs2.262.965120033
tag=TAG20180111T085150
channel ORA_DISK_1: datafile copy complete, elapsed time: 00:00:25
channel ORA_DISK_1: starting datafile copy
input datafile file number=00007 name=+DATAC1/racddb/DATAFILE/users.dbf
output file name=+DATA/RACCDDB_PHX34Z/DATAFILE/users.261.965120057
tag=TAG20180111T085150
channel ORA_DISK_1: datafile copy complete, elapsed time: 00:00:15
channel ORA_DISK_1: starting datafile copy
input datafile file number=00002 name=+DATAC1/racddb/PDBSEED/DATAFILE/system.dbf
output file
name=+DATA/RACCDDB_PHX34Z/4D8C4492AD829E0DE0530C4F800ADE8D/DATAFILE/system.274.9651200
73 tag=TAG20180111T085150
channel ORA_DISK_1: datafile copy complete, elapsed time: 00:00:07
channel ORA_DISK_1: starting datafile copy
input datafile file number=00004 name=+DATAC1/racddb/PDBSEED/DATAFILE/sysaux.dbf
output file
name=+DATA/RACCDDB_PHX34Z/4D8C4492AD829E0DE0530C4F800ADE8D/DATAFILE/sysaux.275.9651200
79 tag=TAG20180111T085150
channel ORA_DISK_1: datafile copy complete, elapsed time: 00:00:07
channel ORA_DISK_1: starting datafile copy
input datafile file number=00008 name=+DATAC1/racddb/tokyo/tokyo_system.dbf
output file
name=+DATA/RACCDDB_PHX34Z/627B610857F86EA9E0537701000A1DB6/DATAFILE/system.276.9651200
87 tag=TAG20180111T085150
channel ORA_DISK_1: datafile copy complete, elapsed time: 00:00:07
channel ORA_DISK_1: starting datafile copy
input datafile file number=00010 name=+DATAC1/racddb/tokyo/tokyo_users01.dbf
output file
name=+DATA/RACCDDB_PHX34Z/627B610857F86EA9E0537701000A1DB6/DATAFILE/users.277.96512009
3 tag=TAG20180111T085150
channel ORA_DISK_1: datafile copy complete, elapsed time: 00:00:01
Finished backup at 11-JAN-18
```



```
sql statement: alter system archive log current
```

```
contents of Memory Script:
```

```
{  
    switch clone datafile all;  
}
```

```
executing Memory Script
```

```
datafile 1 switched to datafile copy
```

```
input datafile copy RECID=1 STAMP=965120096 file  
name=+DATA/RACCDDB_PHX34Z/DATAFILE/system.265.965119957
```

```
datafile 2 switched to datafile copy
```

```
input datafile copy RECID=2 STAMP=965120096 file  
name=+DATA/RACCDDB_PHX34Z/4D8C4492AD829E0DE0530C4F800ADE8D/DATAFILE/system.274.9651200  
73
```

```
datafile 3 switched to datafile copy
```

```
input datafile copy RECID=3 STAMP=965120096 file  
name=+DATA/RACCDDB_PHX34Z/DATAFILE/sysaux.272.965119983
```

```
datafile 4 switched to datafile copy
```

```
input datafile copy RECID=4 STAMP=965120096 file  
name=+DATA/RACCDDB_PHX34Z/4D8C4492AD829E0DE0530C4F800ADE8D/DATAFILE/sysaux.275.9651200  
79
```

```
datafile 5 switched to datafile copy
```

```
input datafile copy RECID=5 STAMP=965120096 file  
name=+DATA/RACCDDB_PHX34Z/DATAFILE/undotbs1.263.965120007
```

```
datafile 6 switched to datafile copy
```

```
input datafile copy RECID=6 STAMP=965120096 file  
name=+DATA/RACCDDB_PHX34Z/DATAFILE/undotbs2.262.965120033
```

```
datafile 7 switched to datafile copy
```

```
input datafile copy RECID=7 STAMP=965120096 file  
name=+DATA/RACCDDB_PHX34Z/DATAFILE/users.261.965120057
```

```
datafile 8 switched to datafile copy
```

```
input datafile copy RECID=8 STAMP=965120096 file  
name=+DATA/RACCDDB_PHX34Z/627B610857F86EA9E0537701000A1DB6/DATAFILE/system.276.9651200  
87
```

```
datafile 9 switched to datafile copy
```

```
input datafile copy RECID=9 STAMP=965120096 file  
name=+DATA/RACCDDB_PHX34Z/627B610857F86EA9E0537701000A1DB6/DATAFILE/sysaux.266.9651199  
11
```

```
datafile 10 switched to datafile copy
```

```
input datafile copy RECID=10 STAMP=965120096 file  
name=+DATA/RACCDDB_PHX34Z/627B610857F86EA9E0537701000A1DB6/DATAFILE/users.277.96512009  
3
```

Appendix B: MAA Best Practice Parameter Settings

The following settings are recommended to follow MAA best practices in order to provide maximum availability and protection of the data. These parameters should be set on both the primary and standby databases.

- ARCHIVELOG enabled
- Flashback database on
- FORCE LOGGING enabled
- Use SPFILE
- Use Data Guard Broker
- COMPATIBLE uses 4 decimals and is the same on both databases
- DB_FILES=1024
- Online Redo Log characteristics
 - At least 1G in size for non-Exadata databases, 4G for Exadata
 - Only multiplexed on normal redundancy storage; single member groups when using high redundancy storage
 - Minimum of 3 online log groups per thread
 - Reside on DATA disk group
- Standby Redo Log characteristics
 - Identical size as online redo logs
 - For RAC, assign SRL groups to a thread
 - Single member only
 - Same number of groups per thread as online redo log groups
 - Reside on DATA disk group
- LOG_BUFFER = 128M for 11.2; 256M for 12.1+
- DB_BLOCK_CHECKING=MEDIUM or FULL *Note: this setting could affect performance and should be enabled only after proper testing of the application.*
- DB_BLOCK_CHECKSUM=TYPICAL
- STANDBY_FILE_MANAGEMENT=AUTO
- DB_LOST_WRITE_PROTECT=TYPICAL
- DB_FLASHBACK_RETENTION_TARGET=minimum of 120
- FAST_START_MTTR_TARGET=300

- USE_LARGE_PAGES=ONLY if hugepages are configured and properly sized on the on-prem system
- CLUSTER_INTERCONNECTS set per gv\$cluster_interconnects **#Exadata Only**
- PARALLEL_THREADS_PER_CPU=1
- DB_CREATE_ONLINE_LOG_DEST_1= DATA disk group
- DB_CREATE_ONLINE_LOG_DEST_n other than 1 should only be set when DATA is not high redundancy
- DB_CREATE_FILE_DEST uses DATA disk group
- DB_RECOVERY_FILE_DEST uses RECO disk group
- Recyclebin is on

Appendix C: Example output for RMAN RESTORE..FROM SERVICE

```
[oracle@host1 ~]$ rman target /
```

```
Recovery Manager: Release 18.0.0.0.0 - Production on Sat Mar 2 16:48:14 2019
Version 18.3.0.0.0
```

```
Copyright (c) 1982, 2018, Oracle and/or its affiliates. All rights reserved.
```

```
connected to target database (not started)
```

```
RMAN> startup nomount
```

```
Oracle instance started
```

```
Total System Global Area 32212253688 bytes
```

```
Fixed Size          12454904 bytes
Variable Size       4160749568 bytes
Database Buffers    27984396288 bytes
Redo Buffers        54652928 bytes
```

```
RMAN> CONFIGURE DEVICE TYPE DISK PARALLELISM 4;
```

```
new RMAN configuration parameters:
```

```
CONFIGURE DEVICE TYPE DISK PARALLELISM 4 BACKUP TYPE TO BACKUPSET;
```

```
new RMAN configuration parameters are successfully stored
```

```
RMAN> restore standby controlfile from service 'victorp';
```

```
Starting restore at 02-MAR-19
```

```
using target database control file instead of recovery catalog
```


```
allocated channel: ORA_DISK_1
```

```
channel ORA_DISK_1: SID=10 instance=victor1 device type=DISK
```

```
channel ORA_DISK_1: starting datafile backup set restore
```

```
channel ORA_DISK_1: using network backup set from service victorp
```

```
channel ORA_DISK_1: restoring control file
```



channel ORA_DISK_1: restore complete, elapsed time: 00:00:20
output file name=+RECO/VICTOR_PHX2W5/CONTROLFILE/current.256.1001775751
Finished restore at 02-MAR-19

RMAN> alter database mount;

released channel: ORA_DISK_1
Statement processed

RMAN> restore database from service 'victorp' section size 2G;

Starting restore at 02-MAR-19
Starting implicit crosscheck backup at 02-MAR-19
allocated channel: ORA_DISK_1
allocated channel: ORA_DISK_2
allocated channel: ORA_DISK_3
allocated channel: ORA_DISK_4
Crosschecked 1 objects
Finished implicit crosscheck backup at 02-MAR-19

Starting implicit crosscheck copy at 02-MAR-19
using channel ORA_DISK_1
using channel ORA_DISK_2
using channel ORA_DISK_3
using channel ORA_DISK_4
Finished implicit crosscheck copy at 02-MAR-19

searching for all files in the recovery area
cataloging files...
cataloging done

List of Cataloged Files

=====


File Name: +RECO/VICTOR_PHX2W5/ARCHIVELOG/2019_03_01/thread_2_seq_1.261.1001776239
File Name: +RECO/VICTOR_PHX2W5/ARCHIVELOG/2019_03_01/thread_2_seq_2.262.1001776491
File Name: +RECO/VICTOR_PHX2W5/ARCHIVELOG/2019_03_01/thread_1_seq_1.263.1001776653
File Name: +RECO/VICTOR_PHX2W5/ARCHIVELOG/2019_03_01/thread_2_seq_3.264.1001776653
File Name: +RECO/VICTOR_PHX2W5/ARCHIVELOG/2019_03_01/thread_1_seq_2.265.1001776679
File Name: +RECO/VICTOR_PHX2W5/ARCHIVELOG/2019_03_01/thread_2_seq_4.266.1001776721
File Name: +RECO/VICTOR_PHX2W5/ARCHIVELOG/2019_03_01/thread_2_seq_5.267.1001776723
File Name: +RECO/VICTOR_PHX2W5/ARCHIVELOG/2019_03_01/thread_1_seq_3.270.1001791759

using channel ORA_DISK_1
using channel ORA_DISK_2
using channel ORA_DISK_3
using channel ORA_DISK_4

channel ORA_DISK_1: starting datafile backup set restore
channel ORA_DISK_1: using network backup set from service victorp
channel ORA_DISK_1: specifying datafile(s) to restore from backup set
channel ORA_DISK_1: restoring datafile 00001 to +DATA/VICTOR_PHX2W5/DATAFILE/system.263.1001869475
channel ORA_DISK_1: restoring section 1 of 1
channel ORA_DISK_2: starting datafile backup set restore
channel ORA_DISK_2: using network backup set from service victorp
channel ORA_DISK_2: specifying datafile(s) to restore from backup set
channel ORA_DISK_2: restoring datafile 00003 to +DATA/VICTOR_PHX2W5/DATAFILE/sysaux.274.1001869485
channel ORA_DISK_2: restoring section 1 of 1
channel ORA_DISK_3: starting datafile backup set restore
channel ORA_DISK_3: using network backup set from service victorp
channel ORA_DISK_3: specifying datafile(s) to restore from backup set



channel ORA_DISK_3: restoring datafile 00004 to
+DATA/VICTOR_PHX2W5/DATAFILE/undotbs1.270.1001869497
channel ORA_DISK_3: restoring section 1 of 1
channel ORA_DISK_4: starting datafile backup set restore
channel ORA_DISK_4: using network backup set from service victorp
channel ORA_DISK_4: specifying datafile(s) to restore from backup set
channel ORA_DISK_4: restoring datafile 00005 to +DATA/VICTOR_PHX2W5/DATAFILE/system.272.1001869507
channel ORA_DISK_4: restoring section 1 of 1
channel ORA_DISK_3: restore complete, elapsed time: 00:00:25
channel ORA_DISK_3: starting datafile backup set restore
channel ORA_DISK_3: using network backup set from service victorp
channel ORA_DISK_3: specifying datafile(s) to restore from backup set
channel ORA_DISK_3: restoring datafile 00006 to +DATA/VICTOR_PHX2W5/DATAFILE/sysaux.271.1001869523
channel ORA_DISK_3: restoring section 1 of 1
channel ORA_DISK_4: restore complete, elapsed time: 00:05:02
channel ORA_DISK_4: starting datafile backup set restore
channel ORA_DISK_4: using network backup set from service victorp
channel ORA_DISK_4: specifying datafile(s) to restore from backup set
channel ORA_DISK_4: restoring datafile 00007 to +DATA/VICTOR_PHX2W5/DATAFILE/users.268.1001869811
channel ORA_DISK_4: restoring section 1 of 1
channel ORA_DISK_4: restore complete, elapsed time: 00:00:19
channel ORA_DISK_4: starting datafile backup set restore
channel ORA_DISK_4: using network backup set from service victorp
channel ORA_DISK_4: specifying datafile(s) to restore from backup set
channel ORA_DISK_4: restoring datafile 00008 to
+DATA/VICTOR_PHX2W5/DATAFILE/undotbs1.266.1001869831
channel ORA_DISK_4: restoring section 1 of 1
channel ORA_DISK_4: restore complete, elapsed time: 00:01:07
channel ORA_DISK_4: starting datafile backup set restore
channel ORA_DISK_4: using network backup set from service victorp
channel ORA_DISK_4: specifying datafile(s) to restore from backup set
channel ORA_DISK_4: restoring datafile 00009 to
+DATA/VICTOR_PHX2W5/DATAFILE/undotbs2.264.1001869899
channel ORA_DISK_4: restoring section 1 of 1
channel ORA_DISK_3: restore complete, elapsed time: 00:06:41
channel ORA_DISK_3: starting datafile backup set restore
channel ORA_DISK_3: using network backup set from service victorp
channel ORA_DISK_3: specifying datafile(s) to restore from backup set
channel ORA_DISK_3: restoring datafile 00010 to +DATA/VICTOR_PHX2W5/DATAFILE/system.265.1001869925
channel ORA_DISK_3: restoring section 1 of 1
channel ORA_DISK_4: restore complete, elapsed time: 00:00:39
channel ORA_DISK_4: starting datafile backup set restore
channel ORA_DISK_4: using network backup set from service victorp
channel ORA_DISK_4: specifying datafile(s) to restore from backup set
channel ORA_DISK_4: restoring datafile 00011 to +DATA/VICTOR_PHX2W5/DATAFILE/sysaux.262.1001869939
channel ORA_DISK_4: restoring section 1 of 1
channel ORA_DISK_3: restore complete, elapsed time: 00:05:10
channel ORA_DISK_3: starting datafile backup set restore
channel ORA_DISK_3: using network backup set from service victorp
channel ORA_DISK_3: specifying datafile(s) to restore from backup set
channel ORA_DISK_3: restoring datafile 00012 to
+DATA/VICTOR_PHX2W5/DATAFILE/undotbs1.261.1001870237
channel ORA_DISK_3: restoring section 1 of 1
channel ORA_DISK_2: restore complete, elapsed time: 00:12:44
channel ORA_DISK_2: starting datafile backup set restore
channel ORA_DISK_2: using network backup set from service victorp
channel ORA_DISK_2: specifying datafile(s) to restore from backup set
channel ORA_DISK_2: restoring datafile 00013 to +DATA/VICTOR_PHX2W5/DATAFILE/undo_2.260.1001870251
channel ORA_DISK_2: restoring section 1 of 1
channel ORA_DISK_3: restore complete, elapsed time: 00:00:27
channel ORA_DISK_3: starting datafile backup set restore



```
channel ORA_DISK_3: using network backup set from service victorp
channel ORA_DISK_3: specifying datafile(s) to restore from backup set
channel ORA_DISK_3: restoring datafile 00014 to +DATA/VICTOR_PHX2W5/DATAFILE/users.280.1001870265
channel ORA_DISK_3: restoring section 1 of 1
channel ORA_DISK_2: restore complete, elapsed time: 00:00:23
channel ORA_DISK_3: restore complete, elapsed time: 00:00:20
channel ORA_DISK_4: restore complete, elapsed time: 00:06:25
channel ORA_DISK_1: restore complete, elapsed time: 00:14:13
Finished restore at 02-MAR-19
```

Appendix D: Example create database JSON file

```
{
  "object": "db",
  "action": "start",
  "operation": "createdb",
  "params": {
    "nodelist": " ",
    "dbname": "sales",
    "db_unique_name": "sales_stby",
    "edition": "EE_EP",
    "version": "18.0.0.0",
    "adminPassword": "W3lCome-#",
    "pdbName": "WEBPDB1",
    "charset": "AL32UTF8",
    "ncharset": "AL16UTF16",
    "backupDestination": "NONE"
  },
  "outputfile": "/home/oracle/createdb/createdb.out",
  "FLAGS": ""
}
```

**Oracle Corporation, World Headquarters**

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

blogs.oracle.com/oracle



facebook.com/oracle



twitter.com/oracle



oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2018, 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.

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. 0719

White Paper Title
July 2019
Author: Kazuhiro Ikeda, Sebastian Solbach

Contributing Authors: Ramachandran Pandrapattahil



| Oracle is committed to developing practices and products that help protect the environment.