Initialization Parameters

It provides complete initialization parameter information, including how to update initialization parameters by issuing the ALTER SYSTEM SET statement (for example, ALTER SYSTEM SET LOG_ARCHIVE_TRACE) or by editing the initialization parameter files.

See the Oracle operating system-specific documentation for more information about setting initialization parameters.

Table-1: Initialization Parameters for Instances in an Oracle Data Guard Configuration

Parameter	Applicable To	Notes and Recommendations
COMPATIBLE = *release_number*	Primary Logical Standby Physical Standby Snapshot Standby	Specify the same value on the primary and standby databases if you expect to do a switchover. If the values differ, redo transport services may be unable to transmit redo data from the primary database to the standby databases
CONTROL_FILE_RECORD_KEEP_TIME = *number_of_days*	Primary Logical Standby Physical Standby Snapshot Standby	Optional. Use this parameter to avoid overwriting a reusable record in the control file (that contains needed information such as an archived redo log file) for the specified number of days (from 0 to 365).
CONTROL_FILES = '*control_file_name*', '*control_file_name*', '...'	Primary Logical Standby Physical Standby Snapshot Standby	Required. Specify the path name and filename for one or more control files. The control files must already exist on the database. Oracle recommends using 2 control files. If another copy of the current control file is available, then an instance can be easily restarted after copying the good control file to the location of the bad control file.
DB_FILE_NAME_CONVERT = '*location_of_primary_database_datafile*','*location_of_standby_database_datafile*'	Physical Standby	This parameter must specify paired strings. The first string is a sequence of characters to be looked for in a primary

Parameter	Applicable To	Notes and Recommendations
	Snapshot Standby	database filename. If that sequence of characters is matched, it is replaced by the second string to construct the standby database filename. You can specify multiple pairs of filenames.
DB_UNIQUE_NAME = *Unique name for the database*	Primary Logical Standby Physical Standby Snapshot Standby	Recommended, but required if you specify the LOG_ARCHIVE_CONFIG parameter. Specifies a unique name for this database. This name does not change even if the primary and standby databases reverse roles. The DB_UNIQUE_NAME parameter defaults to the value of the DB_NAME parameter.
FAL_CLIENT = *Oracle_Net_service_name*	Physical Standby Snapshot Standby	This parameter is no longer required. If it is not set, the fetch archive log (FAL) server obtains the client's network address from the LOG_ARCHIVE_DEST_*n* parameter that corresponds to the client's DB_UNIQUE_NAME.
FAL_SERVER = *Oracle_Net_service_name*	Physical Standby Snapshot Standby	Specifies one or more Oracle Net service names for the databases from which this standby database can fetch (request) missing archived redo log files.
INSTANCE_NAME	Primary Logical Standby Physical Standby Snapshot Standby	Optional. If this parameter is defined and the primary and standby databases reside on the same host, specify a different name for the standby database than you specify for the primary database.
LOG_ARCHIVE_CONFIG ='DG_CONFIG=(*db_unique_name, db_unique_name, ...*)'	Primary Logical Standby	Highly recommended. The DG_CONFIG attribute of this parameter must be explicitly set on each database in an Oracle Data Guard

Parameter	Applicable To	Notes and Recommendations
	Physical Standby Snapshot Standby	configuration to enable full Oracle Data Guard functionality. Set DG_CONFIG to a text string that contains the DB_UNIQUE_NAME of each database in the configuration, with each name in this list separated by a comma.
LOG_ARCHIVE_DEST_*n* = {LOCATION=*path_name* | SERVICE=*service_name, attribute, attribute, ...*}	Primary Logical Standby Physical Standby Snapshot Standby	Required. Define up to thirty (where *n* = 1, 2, 3, ... 31) destinations, each of which must specify either the LOCATION or SERVICE attribute. Specify a corresponding LOG_ARCHIVE_DEST_STATE_*n* parameter for every LOG_ARCHIVE_DEST_*n* parameter.
LOG_ARCHIVE_DEST_STATE_*n* = {ENABLE|DEFER|ALTERNATE}	Primary Logical Standby Physical Standby Snapshot Standby	Required. Specify a LOG_ARCHIVE_DEST_STATE_*n* parameter to enable or disable redo transport services to transmit redo data to the specified (or to an alternate) destination. Define a LOG_ARCHIVE_DEST_STATE_*n* parameter for every LOG_ARCHIVE_DEST_*n* parameter..
LOG_ARCHIVE_FORMAT=*log%d_%t_%s_%r.arc*	Primary Logical Standby Physical Standby Snapshot Standby	The LOG_ARCHIVE_FORMAT and LOG_ARCHIVE_DEST_n parameters are concatenated together to generate fully qualified archived redo log filenames on a database.
LOG_ARCHIVE_MAX_PROCESSES =*integer*	Primary Logical Standby Physical Standby Snapshot Standby	Optional. Specify the number (from 1 to 30) of archiver (ARC*n*) processes you want Oracle software to invoke initially. The default value is 4.

Parameter	Applicable To	Notes and Recommendations
LOG_ARCHIVE_MIN_SUCCEED_DEST	Primary Logical Standby Physical Standby Snapshot Standby	Optional. This parameter specifies the number of local or remote MANDATORY destinations, or local OPTIONAL destinations, that a logfile group must be archived to before it can be re-used.
LOG_ARCHIVE_TRACE=*integer*	Primary Logical Standby Physical Standby Snapshot Standby	Optional. Set this parameter to trace the transmission of redo data to the standby site.
LOG_FILE_NAME_CONVERT = '*location_of_primary_database_redo_logs*','*location_of_standby_database_redo_logs*'	Logical Standby Physical Standby Snapshot Standby	This parameter converts the path names of the primary database online redo log file to path names on the standby database.
REMOTE_LOGIN_PASSWORDFILE = {EXCLUSIVE|SHARED}	Primary Logical Standby Physical Standby Snapshot Standby	Optional if operating system authentication is used for administrative users and SSL is used for redo transport authentication. Otherwise, this parameter must be set to EXCLUSIVE or SHARED on every database in an Oracle Data Guard configuration.
SHARED_POOL_SIZE_= bytes	Primary Logical Standby Physical Standby Snapshot Standby	Optional. Use to specify the system global area (SGA) to stage the information read from the online redo log files. The more SGA that is available, the more information that can be staged.

Parameter	Applicable To	Notes and Recommendations
`STANDBY_FILE_MANAGEMENT = {AUTO | MANUAL}`	Primary Physical Standby Snapshot Standby	Set this parameter to AUTO so that when data files are added to or dropped from the primary database, corresponding changes are made in the standby database without manual intervention. If the directory structures on the primary and standby databases are different, you must also set the DB_FILE_NAME_CONVERT in itialization parameter to convert the filenames of one or more sets of data files on the primary database to filenames on the (physical) standby database.

LOG_ARCHIVE_DEST_n Parameter Attributes

This is a list of the attributes for the LOG_ARCHIVE_DEST_*n* initialization parameter, (where *n* is an integer between 1 and 31).

- AFFIRM and NOAFFIRM
- COMPRESSION
- DB_UNIQUE_NAME
- DELAY
- SYNC and ASYNC (SYNC is not supported or LOG_ARCHIVE_DEST_11 through LOG_ARCHIVE_DEST_31)
- TEMPLATE
- VALID_FOR

- Each database in an Oracle Data Guard configuration typically has one destination with the LOCATION attribute for the archival of the online and standby redo logs and one destination with the REMOTE attribute for every other database in the configuration.
- If configured, each LOG_ARCHIVE_DEST_1 through LOG_ARCHIVE_DEST_10 destination must contain either a LOCATION or SERVICE attribute to specify a local disk directory or a remotely accessed database, respectively.
  Each LOG_ARCHIVE_DEST_11 through LOG_ARCHIVE_DEST_31 destination must contain a SERVICE attribute.
  All other attributes are optional.
- LOG_ARCHIVE_DEST_11 through LOG_ARCHIVE_DEST_31 can only be used when the COMPATIBLE initialization parameter is set to 11.2.0.0 or later.

SYNC and ASYNC

The SYNC and ASYNC attributes specify whether the synchronous (SYNC) or asynchronous (ASYNC) redo transport mode is to be used.

Category	SYNC	ASYNC
Data type	Keyword	Keyword
Valid values	Not applicable	Not applicable
Default value	Not applicable	None
Requires attributes	None	None
Conflicts with attributes	ASYNC, LOCATION	SYNC, LOCATION
Corresponds to	TRANSMIT_MODE column of the V$ARCHIVE_DEST view	TRANSMIT_MODE column of the V$ARCHIVE_DEST view

Usage Notes

- The LOG_ARCHIVE_DEST_11 through LOG_ARCHIVE_DEST_31 parameters do not support the SYNC attribute.
- The redo data generated by a transaction must have been received by every enabled destination that has the SYNC attribute before that transaction can commit.
- On primary databases and logical standbys, destinations 1 through 10 default to ASYNC (real-time cascading).
  On physical standbys, snapshot standbys, and far sync instances, destinations 1 through 10 default to ARCH transport mode. (Note that the ARCH attribute is deprecated; the use of ARCH in this situation simply indicates non-real-time cascading.)
  Destinations 11 through 31 always default to ASYNC.

Example

The following example shows the SYNC attribute with the LOG_ARCHIVE_DEST_*n* parameter.

```
LOG_ARCHIVE_DEST_3='SERVICE=stby1 SYNC'

LOG_ARCHIVE_DEST_STATE_3=ENABLE
```

AFFIRM and NOAFFIRM

The AFFIRM and NOAFFIRM attributes control whether a redo transport destination acknowledges received redo data before or after writing it to the standby redo log.

Definitions of each option are as follows:

- AFFIRM—specifies that a redo transport destination acknowledges received redo data after writing it to the standby redo log.
- NOAFFIRM—specifies that a redo transport destination acknowledges received redo data before writing it to the standby redo log.

Category	AFFIRM	NOAFFIRM
Data type	Keyword	Keyword
Valid values	Not applicable	Not applicable
Default Value	Not applicable	Not applicable
Requires attributes	SERVICE	SERVICE
Conflicts with attributes	NOAFFIRM	AFFIRM
Corresponds to	AFFIRM column of the V$ARCHIVE_DEST view	AFFIRM column of the V$ARCHIVE_DEST view

Usage Notes

- If neither the AFFIRM nor the NOAFFIRM attribute is specified, then the default is AFFIRM when the SYNC attribute is specified and NOAFFIRM when the ASYNC attribute is specified.
- Specification of the AFFIRM attribute without the SYNC attribute is deprecated and will not be supported in future releases.

```
LOG_ARCHIVE_DEST_3='SERVICE=stby1 SYNC AFFIRM'
LOG_ARCHIVE_DEST_STATE_3=ENABLE
```

COMPRESSION

The COMPRESSION attribute is used to specify whether redo data is compressed before transmission to a redo transport destination.

```
LOG_ARCHIVE_DEST_3='SERVICE=denver SYNC COMPRESSION=ENABLE'
LOG_ARCHIVE_DEST_STATE_3=ENABLE
```

DB_UNIQUE_NAME

The DB_UNIQUE_NAME attribute specifies a unique name for the database at this destination.

- This attribute is optional if:
  - The LOG_ARCHIVE_CONFIG=DG_CONFIG initialization parameter is not specified.
  - This is a local destination (specified with the LOCATION attribute).
- This attributes is required if the LOG_ARCHIVE_CONFIG=DG_CONFIG initialization parameter is specified and if this is a remote destination (specified with the SERVICE attribute).
- Use the DB_UNIQUE_NAME attribute to clearly identify the relationship between a primary and standby databases. This attribute is particularly helpful if there are multiple standby databases in the Oracle Data Guard configuration.
- The name specified by the DB_UNIQUE_NAME must match one of the DB_UNIQUE_NAME values in the DG_CONFIG list. Redo transport services validate that the DB_UNIQUE_NAME attribute of the database at the specified destination matches the DB_UNIQUE_NAME attribute or the connection to that destination is refused.
- The name specified by the DB_UNIQUE_NAME attribute must match the name specified by the DB_UNIQUE_NAME initialization parameter of the database identified by the destination.

```
DB_UNIQUE_NAME=boston
LOG_ARCHIVE_CONFIG='DG_CONFIG=(chicago,boston,denver)'
LOG_ARCHIVE_DEST_1='LOCATION=/arch1/
  VALID_FOR=(ALL_LOGFILES,ALL_ROLES)
  DB_UNIQUE_NAME=boston'
LOG_ARCHIVE_DEST_2='SERVICE=Sales_DR
  VALID_FOR=(ONLINE_LOGFILES,PRIMARY_ROLE)
  DB_UNIQUE_NAME=chicago'
```

DELAY

The DELAY attribute specifies a minimum time lag between when redo data from the primary site is archived on a standby site and when the archived redo log file is applied to the standby database or any standbys cascaded from it.

- The DELAY attribute is optional. By default there is no delay.
- The DELAY attribute indicates the archived redo log files at the standby destination are not available for recovery until the specified time interval has expired. The time interval is expressed in minutes, and it starts when the redo data is successfully transmitted to, and archived at, the standby site.
- The DELAY attribute may be used to protect a standby database from corrupted or erroneous primary data. However, there is a tradeoff because during failover it takes more time to apply all of the redo up to the point of corruption.
- The DELAY attribute does not affect the transmittal of redo data to a standby destination.
- If you have real-time apply enabled, then any delay that you set is ignored.

- Changes to the DELAY attribute take effect the next time redo data is archived (after a log switch). In-progress archiving is not affected.
- You can override the specified delay interval at the standby site, as follows:
    - For a physical standby database:

```
SQL> ALTER DATABASE RECOVER MANAGED STANDBY DATABASE NODELAY;
```

For a logical standby database:

```
SQL> ALTER DATABASE START LOGICAL STANDBY APPLY NODELAY;
```

```
LOG_ARCHIVE_DEST_1='LOCATION=/arch/dest MANDATORY'
LOG_ARCHIVE_DEST_STATE_1=ENABLE
LOG_ARCHIVE_DEST_2='SERVICE=stbyB SYNC AFFIRM'
LOG_ARCHIVE_DEST_STATE_2=ENABLE
LOG_ARCHIVE_DEST_3='SERVICE=stbyC DELAY=120'
LOG_ARCHIVE_DEST_STATE_3=ENABLE
```

TEMPLATE

The TEMPLATE attribute defines a directory specification and format template for names of archived redo log files at the destination.

Usage Notes

- The TEMPLATE attribute is optional. If TEMPLATE is not specified, archived redo logs are named using the value of the LOG_ARCHIVE_FORMAT initialization parameter.
- The TEMPLATE attribute overrides the LOG_ARCHIVE_FORMAT initialization parameter setting at the remote archival destination.
- The TEMPLATE attribute is valid only with remote destinations (specified with the SERVICE attribute).
- The value you specify for filename_template must contain the %s, %t, and %r directives described in below table:

Directive	Description
%t	Substitute the instance thread number.
%T	Substitute the instance thread number, zero filled.
%s	Substitute the log file sequence number.
%S	Substitute the log file sequence number, zero filled.

Directive	Description
%r	Substitute the resetlogs ID.
%R	Substitute the resetlogs ID, zero filled.

VALID_FOR

The VALID_FOR attribute specifies whether redo data gets written to a destination.

The following factors are considered:

- Whether the database is currently running in the primary or the standby role
- Whether online redo log files, standby redo log files, or both are currently being archived on the database at this destination

Category	VALID_FOR=(redo_log_type, database_role)
Data Type	String value
Valid values	Not applicable
Default Value	VALID_FOR=(ALL_LOGFILES, ALL_ROLES)
Requires attributes	None
Conflicts with attributes	None
Corresponds to	VALID_NOW, VALID_TYPE, and VALID_ROLE columns in the V$ARCHIVE_DEST view

Usage Notes

- The VALID_FOR attribute is optional. However, Oracle recommends that the VALID_FOR attribute be specified for each redo transport destination at each database in an Oracle Data Guard configuration so that redo transport continues after a role transition to any standby database in the configuration.

- To configure these factors for each LOG_ARCHIVE_DEST_*n* destination, you specify this attribute with a pair of keywords: VALID_FOR=(*redo_log_type*,*database_role*):

    - The *redo_log_type* keyword identifies the destination as valid for archiving one of the following:

    - ONLINE_LOGFILE—This destination is valid only when archiving online redo log files.
    - STANDBY_LOGFILE—This destination is valid only when archiving standby redo log files.
    - ALL_LOGFILES— This destination is valid when archiving either online redo log files or standby redo log files.

  - The *database_role* keyword identifies the role in which this destination is valid for archiving:
    - PRIMARY_ROLE—This destination is valid only when the database is running in the primary role.
    - STANDBY_ROLE—This destination is valid only when the database is running in the standby role.
    - ALL_ROLES—This destination is valid when the database is running in either the primary or the standby role.

- If you do not specify the VALID_FOR attribute for a destination, by default, archiving online redo log files and standby redo log files is enabled at the destination, regardless of whether the database is running in the primary or the standby role. This default behavior is equivalent to setting the (ALL_LOGFILES,ALL_ROLES) keyword pair on the VALID_FOR attribute.
- The VALID_FOR attribute enables you to use the same initialization parameter file for both the primary and standby roles.

Example

The following example shows the default VALID_FOR keyword pair:

```
LOG_ARCHIVE_DEST_1='LOCATION=/disk1/oracle/oradata VALID_FOR=(ALL_LOGFILES,
ALL_ROLES)'
```