

Oracle FLEXCUBE Products Database Practices

Db 19c

Release 14.4.0.0.0

August 2020



Table of Contents

| | |
|--|------------|
| 1. INTRODUCTION | 1-1 |
| 2. DATABASE INITIALIZATION PARAMETERS | 2-1 |
| 2.1 ALLOW_LEVEL_WITHOUT_CONNECT_BY | 2-1 |
| 2.2 LM_RES_HASH_BUCKET | 2-1 |
| 2.3 CURSOR_SHARING | 2-1 |
| 2.4 DB_CACHE_ADVICE | 2-2 |
| 2.5 FAST_START_MTTR_TARGET | 2-2 |
| 2.6 JOB_QUEUE_PROCESSES | 2-3 |
| 2.7 LOG_BUFFER | 2-3 |
| 2.8 MEMORY_TARGET/MEMORY_MAX_TARGET | 2-3 |
| 2.9 NLS_DATE_FORMAT | 2-4 |
| 2.10 OPEN_CURSORS | 2-4 |
| 2.11 OPTIMIZER_DYNAMIC_SAMPLING | 2-5 |
| 2.12 OPTIMIZER_INDEX_CACHING | 2-5 |
| 2.13 OPTIMIZER_INDEX_COST_ADJ | 2-6 |
| 2.14 PARALLEL_MAX_SERVERS | 2-6 |
| 2.15 PGA_AGGREGATE_LIMIT | 2-7 |
| 2.16 PLSQL_CODE_TYPE | 2-7 |
| 2.17 PROCESSES | 2-7 |
| 2.18 REMOTE_DEPENDENCIES_MODE | 2-7 |
| 2.19 SESSION_CACHED_CURSORS | 2-8 |
| 2.20 SKIP_UNUSABLE_INDEXES | 2-8 |
| 2.21 UNDO_RETENTION | 2-9 |
| 3. REDO LOG FILES | 3-1 |
| 4. PLSQL OPTIMIZER LEVEL | 4-1 |
| 5. STATISTICS COLLECTION FOR FLEXCUBE SCHEMA (RECOMMENDED METHOD) | 5-1 |
| 5.1 CUSTOMIZING DEFAULT STATISTICS COLLECTION SCHEDULE | 5-1 |
| 5.2 CUSTOMIZING STATISTICS GATHERING FOR FLEXCUBE | 5-3 |
| 5.2.1 <i>Statistics Histograms</i> | 5-3 |
| 6. FLEXCUBE DATABASE STORAGE RECOMMENDATIONS | 6-1 |
| 6.1 KEY BENEFITS OF ASM | 6-1 |
| 7. FLEXCUBE DATABASE BACKUP RECOMMENDATIONS | 7-1 |
| 7.1 RMAN Vs CONVENTIONAL BACKUP | 7-1 |
| 7.2 BENEFITS OF USING RMAN | 7-2 |
| 7.3 BACKUP STRATEGY RECOMMENDATION | 7-3 |
| 8. KNOWN ISSUES | 8-1 |
| 9. APPENDIX | 9-1 |
| 9.1 SCRIPT TO CHECK HISTOGRAMS ON FLEXCUBE SCHEMA | 9-1 |
| 9.2 SCRIPT TO REMOVE HISTOGRAMS ON FLEXCUBE SCHEMA | 9-1 |

1. Introduction

This document details the database setting which is common across all the FLEXCUBE products which include FLEXCUBE Universal Banking Solution, Oracle Banking Payments, Oracle FLEXCUBE Enterprise Limits and Collateral Management, Oracle FLEXCUBE Investor Servicing, Oracle Banking Corporate Lending, Oracle Banking Trade Finance Process Management and Oracle Banking Treasury Management.

For product related database setting, please refer to the product specific database best practice document as well.

2. Database Initialization Parameters

Oracle standard database initialization parameters have been derived after performing the required benchmark tests (Performance Load tests). We recommend installing FLEXCUBE in Pluggable database and few of these parameters need to be set at PDB level.

Note: Since some of the initialization parameters values are specific to customer volume, parameters should be derived using **FLEXCUBE-Disk-Layouts-initparams-19c.xlsx** excel sheet baselined along with this document.

Following are the Parameters with the details and its relevance to FLEXCUBE:

2.1 ALLOW LEVEL WITHOUT CONNECT BY

Recommended Value: TRUE

This parameter is set to avoid following error,

- After Upgrading To Oracle 10g, Getting ORA-01788 When Running A Query That Includes The LEVEL Pseudo Column [ID 455953.1]

2.2 LM RES HASH BUCKET

Recommended Value: 8192

To avoid "LATCH: GES RESOURCE HASH LIST" waits in 19c RAC environment, we need to set this parameter.

2.3 CURSOR SHARING

Determines what kind of SQL statements can share the same cursors.

| Property | Description |
|-------------------|-------------|
| Parameter type | String |
| Default value | EXACT |
| Recommended Value | Force |

Oracle FLEXCUBE relevance

Some of the FLEXCUBE sql statements are generated dynamically. So they contain literal values in the WHERE clause conditions. This results in large numbers of nearly identical statements with separate parse trees in Oracle's library cache, which can slow performance and cause latch problems.

By setting cursor_sharing to FORCE database convert literals to bind variables before parsing the statement.

2.4 DB CACHE ADVICE

This enables or disables statistics gathering used for predicting behavior with different cache sizes through the V\$DB_CACHE_ADVICE performance view.

| Property | Description |
|-------------------|---|
| Parameter type | String |
| Syntax | DB_CACHE_ADVICE = { ON READY OFF } |
| Default value | If STATISTICS_LEVEL is set to TYPICAL / ALL, then ON If STATISTICS_LEVEL is set to BASIC, then OFF |
| Recommended Value | OFF (Should be ON while Performance Monitoring) |

Oracle FLEXCUBE relevance

Turning ON advisory will have an extra overhead. Please note it should be ON, only during performance monitoring.

2.5 FAST START MTTR TARGET

This enables you to specify the number of seconds the database takes to perform crash recovery of a single instance. When specified, FAST_START_MTTR_TARGET is overridden by LOG_CHECKPOINT_INTERVAL.

| Property | Description |
|--------------------|-------------------|
| Parameter type | Integer |
| Default value | 0 |
| Range of values | 0 to 3600 seconds |
| Recommended Values | 300 |

Oracle FLEXCUBE relevance

If FAST_START_MTTR_TARGET is not set to 300 then run time performance for write/redo generation intensive workloads will not be optimized. This will reduce checkpoint writes from DBWR processes, making more room for LGWR IO. To optimize run time performance for write/redo generation intensive workloads, increase the FAST_START_MTTR_TARGET initialization parameter to 300.

2.6 JOB_QUEUE_PROCESSES

This specifies the maximum number of processes that can be created for the execution of jobs. It specifies the number of job queue processes per instance (J000, J999).

| Property | Description |
|--------------------|--|
| Parameter type | Integer |
| Default value | 4000 |
| Range of values | 0 to 1000 |
| Recommended Values | Refer FLEXCUBE-Disk-Layouts-initparams-19c.xlsx |

Oracle FLEXCUBE relevance

This parameter has to be set with respect to the maximum number of scheduler jobs. To arrive at the right value, refer FLEXCUBE-Disk-Layouts-initparams-19c.xlsx excel.

2.7 LOG_BUFFER

Recommended Value: Refer **FLEXCUBE-Disk-Layouts-initparams-19c.xlsx**

Oracle FLEXCUBE relevance

The default log buffer size is too small as FLEXCUBE performs heavy DML during batch processing.

2.8 MEMORY_TARGET/MEMORY_MAX_TARGET

Recommended Value: Refer **FLEXCUBE-Disk-Layouts-initparams-19c.xlsx**.

For linux systems, make sure that the value of operating system /dev/shm mount is set to appropriate value to accommodate memory_Target.

2.9 NLS_DATE_FORMAT

This specifies the default date format to use with the TO_CHAR and TO_DATE functions.

| Property | Description |
|--------------------|----------------------------|
| Parameter type | String |
| Syntax | NLS_DATE_FORMAT = "format" |
| Default value | Derived from NLS_TERRITORY |
| Recommended Values | DD-MON-RRRR |

Oracle FLEXCUBE relevance

FLEXCUBE standard date format.

2.10 OPEN_CURSORS

This specifies the maximum number of open cursors (handles to private SQL areas) a session can have at once. You can use this parameter to prevent a session from opening an excessive number of cursors.

| Property | Description |
|--------------------|---------------------------|
| Parameter type | Integer |
| Default value | 50 |
| Modifiable | ALTER SYSTEM |
| Range of values | 1 to 4294967295 (4 GB -1) |
| Recommended Values | 5000 |

Oracle FLEXCUBE relevance

It is important to set the value of OPEN_CURSORS high enough to prevent FLEXCUBE application from running out of open cursors (ORA-01000: maximum open cursors exceeded).

2.11 OPTIMIZER_DYNAMIC_SAMPLING

This controls the level of dynamic sampling performed by the optimizer.

| Property | Description |
|--------------------|--|
| Parameter type | Integer |
| Default value | If OPTIMIZER_FEATURES_ENABLE is set to 10.0.0 or higher, then 2 If OPTIMIZER_FEATURES_ENABLE is set to 9.2.0, then 1 If OPTIMIZER_FEATURES_ENABLE is set to 9.0.1 or lower, then 0 |
| Recommended Values | 1 |
| Range of values | 0 to 10 |

Oracle FLEXCUBE relevance

Dynamic Sampling is a method of gathering additional statistics during optimization by recursively sampling statements. When dynamic sampling is enabled, queries are recursively generated by Oracle to test various selectivity based upon real values in order to improve their accuracy. This can result in the production of better explain plans.

Value 1 Sample all tables that have not been analyzed that meet certain criteria.

2.12 OPTIMIZER_INDEX_CACHING

This lets you adjust the behavior of cost-based optimization to favor nested loops joins and IN-list iterators.

| Property | Description |
|--------------------|-------------|
| Parameter type | Integer |
| Default value | 0 |
| Recommended Values | 90 |
| Range of values | 0 to 100 |

Oracle FLEXCUBE relevance

The cost of executing an index using IN-list iterators or of executing nested loops join when an index is used to access the inner table depends on the caching of that index in the buffer cache. FLEXCUBE favors nested loop joins by setting optimizer_index_caching to 90.

2.13 OPTIMIZER INDEX COST ADJ

This lets you tune optimizer behavior for access path selection to be more or less index friendly - that is, to make the optimizer more or less prone to selecting an index access path over a full table scan.

| Property | Description |
|--------------------|-------------|
| Parameter type | Integer |
| Default value | 100 |
| Recommended Values | 50 |
| Range of values | 1 to 10000 |

Oracle FLEXCUBE relevance

FLEXCUBE favors index read over full table scan as it is very useful when optimizer favors to give a lower cost to index scans over full-table scans.

2.14 PARALLEL MAX SERVERS

This specifies the maximum number of parallel execution processes and parallel recovery processes for an instance. As demand increases, Oracle Database increases the number of processes from the number created at instance startup up to this value.

| Property | Description |
|---------------------------|--|
| Parameter type | Integer |
| Default value | Derived from the values of CPU_COUNT, PARALLEL_THREADS_PER_CPU, and PGA_AGGREGATE_TARGET |
| Recommended Values | Refer FLEXCUBE-Disk-Layouts-initparams-19c.xlsx |
| Range of values | 0 to 3600 |
| Real Application Clusters | Multiple instances can have different values |

Oracle FLEXCUBE relevance

To arrive right value, refer **FLEXCUBE-Disk-Layouts-initparams-19c.xlsx** excel.

2.15 PGA AGGREGATE LIMIT

Recommended Value: 0

Oracle FLEXCUBE relevance

Setting this parameter limits the pga consumed by the instance, hence might cause failure to few of the running processes.

2.16 PLSQL CODE TYPE

This specifies the compilation mode of the PL/SQL units.

| Property | Description |
|--------------------|---------------------|
| Parameter type | String |
| Default value | INTERPRETED |
| Recommended values | NATIVE |
| Range of values | INTERPRETED, NATIVE |

Oracle FLEXCUBE relevance

The PL SQL interpreter overhead will be minimal when set to NATIVE.

2.17 PROCESSES

This specifies the maximum number of operating system user processes that can simultaneously connect to Oracle. Its value should allow for all background processes such as locks, job queue processes, and parallel execution processes.

| Property | Description |
|--------------------|--|
| Parameter type | Integer |
| Default value | 100 |
| Range of values | 6 to operating system dependent |
| Recommended values | Refer FLEXCUBE-Disk-Layouts-initparams-19c.xlsx |

Oracle FLEXCUBE relevance

This parameter can set be set with respect to maximum no of sessions connected to DB.

2.18 REMOTE DEPENDENCIES MODE

Specifies how Oracle should handle dependencies upon remote PL/SQL stored procedures.

| Property | Description |
|--------------------|--|
| Parameter type | String |
| Syntax | REMOTE_DEPENDENCIES_MODE = { TIMESTAMP SIGNATURE } |
| Default value | TIMESTAMP |
| Recommended values | SIGNATURE |

Oracle FLEXCUBE relevance

Oracle allows the procedure to execute as long as the signatures are considered safe. This setting allows client PL/SQL applications to be run without recompilation.

2.19 SESSION CACHED CURSORS

Specifies the number of session cursors to cache. Repeated parse calls of the same SQL statement cause the session cursor for that statement to be moved into the session cursor cache. Subsequent parse calls will find the cursor in the cache and do not need to reopen the cursor. Oracle uses a least recently used algorithm to remove entries in the session cursor cache to make room for new entries when needed.

| Property | Description |
|--------------------|---------------------------------|
| Parameter type | Integer |
| Default value | 50 |
| Recommended values | 400 |
| Range of values | 0 to operating system-dependent |

Oracle FLEXCUBE relevance

This helps to cache the cursor thus avoid parsing of the cursor which heavy CPU intensive particularly in batch.

2.20 SKIP UNUSABLE INDEXES

Enables or disables the use and reporting of tables with unusable indexes or index partitions.

| Property | Description |
|--------------------|-------------|
| Parameter type | Boolean |
| Default value | True |
| Recommended values | FALSE |

| | |
|-----------------|--------------|
| Range of values | true / false |
|-----------------|--------------|

Oracle FLEXCUBE relevance

TRUE enables error reporting of indexes marked UNUSABLE. This setting does not allow inserts, deletes, and updates on tables with unusable indexes or index partitions. IT is set to false because FLEXCUBE application should throw error if any of the indexes become UNUSABLE.

2.21 UNDO RETENTION

This specifies (in seconds) the low threshold value of undo retention. For AUTOEXTEND undo tablespaces, the system retains undo for at least the time specified in this parameter, and automatically tunes the undo retention period to satisfy the undo requirements of the queries. For fixed- size undo tablespaces, the system automatically tunes for the maximum possible undo retention period, based on undo tablespace size and usage history, and ignores UNDO_RETENTION unless retention guarantee is enabled.

The UNDO_RETENTION parameter can only be honored if the current undo tablespace has enough space. If an active transaction requires undo space and the undo tablespace does not have available space, then the system starts reusing unexpired undo space. This action can potentially cause some queries to fail with a "snapshot too old" message.

| Property | Description |
|--------------------|--------------|
| Parameter type | Integer |
| Default value | 900 |
| Range of values | 0 to 231 – 1 |
| Recommended values | 1800 |

Oracle FLEXCUBE relevance

Increased value along with automatic undo management helps to avoid “snapshot too old error”.

3. Redo Log Files

The default number of redo log files groups and its size is inadequate to run FLEXCUBE. Hence, the recommended are:

- 6 redo log groups
- Redo log file size
 - 1 GB each for the DB size up to 1 TB
 - 2 GB each for DB size more than 1 TB

4. PLSQL Optimizer Level

The `plsql_optimize_level` value for all the pl/sql units should be same which would be the value set in `plsql_optimize_level` init parameter.

Following sql gives the PLSQL optimizer level for FLEXCUBE schema plsql units:

```
Select PLSQL_OPTIMIZE_LEVEL, type, count(*) "Count" from
user_plsql_object_settings group by PLSQL_OPTIMIZE_LEVEL, type;
```

PLSQL_OPTIMIZE_LEVEL for all the objects should be same which should be value set in `plsql_optimize_level` init parameter. If there is a difference then the objects should be recompiled. This can be done using `dbms_utility.compile_schema` procedure.

E.g.: - `exec dbms_utility.compile_schema('FCCBM2')`

Here, 'FCCBM2' refers to the FLEXCUBE schema.

Note: The 'dbms_utility.compile_schema' procedure invalidates and recompiles all the plsql units.

5. Statistics Collection for FLEXCUBE Schema (Recommended Method)

Oracle 19c provides a default scheduled job to collect statistics for the entire database and is default scheduled to run every night. Given that the FLEXCUBE batch as well runs in the night it is critical that the statistics gathering is not run during the batch.

It is recommended to use the default database scheduled job that is shipped with Oracle Database to collect statistics for FLEXCUBE Schema.

Statistics collection recommendation is specific to FLEXCUBE schema. Other available statistics like System statistics, Fixed table statistics, dictionary statistics, etc. are not part of this recommendation and are required to be executed on need basis.

Note: This document assumes that there is no other tool or a program is scheduled to collect statistics for the Database.

5.1 Customizing Default Statistics Collection Schedule

The Default Scheduler is to be customized for the following:

- Ensure that the default statistics gathering program is configured and Running.

```
SELECT STATUS

FROM DBA_AUTOTASK_CLIENT

WHERE CLIENT_NAME='auto optimizer stats collection';
```

Should return - ENABLED

- Ensure that the default statistics gathering program is configured to run only on weekends.

```
/* Start of Script – Script to be executed as SYS*/

BEGIN

DBMS_AUTO_TASK_ADMIN.ENABLE (

    CLIENT_NAME => 'auto optimizer stats collection',

    OPERATION   => NULL,

    WINDOW_NAME => 'SATURDAY_WINDOW');

DBMS_AUTO_TASK_ADMIN.ENABLE (

    CLIENT_NAME => 'auto optimizer stats collection',

    OPERATION   => NULL,

    WINDOW_NAME => 'SUNDAY_WINDOW');
```

```
END;
```

```
/
```

```
/* End of Script */
```

- Default schedule is daily. So disable the daily schedules for optimizer statistics.

```
/* Start of Script – Script to be executed as SYS*/
```

```
BEGIN
```

```
DBMS_AUTO_TASK_ADMIN.DISABLE(
```

```
    CLIENT_NAME => 'auto optimizer stats collection',
```

```
    OPERATION    => NULL,
```

```
    WINDOW_NAME => 'MONDAY_WINDOW');
```

```
DBMS_AUTO_TASK_ADMIN.DISABLE(
```

```
    CLIENT_NAME => 'auto optimizer stats collection',
```

```
    OPERATION    => NULL,
```

```
    WINDOW_NAME => 'TUESDAY_WINDOW');
```

```
DBMS_AUTO_TASK_ADMIN.DISABLE(
```

```
    CLIENT_NAME => 'auto optimizer stats collection',
```

```
    OPERATION    => NULL,
```

```
    WINDOW_NAME => 'WEDNESDAY_WINDOW');
```

```
DBMS_AUTO_TASK_ADMIN.DISABLE(
```

```
    CLIENT_NAME => 'auto optimizer stats collection',
```

```
    OPERATION    => NULL,
```

```
    WINDOW_NAME => 'THURSDAY_WINDOW');
```

```
DBMS_AUTO_TASK_ADMIN.DISABLE(
```

```
    CLIENT_NAME => 'auto optimizer stats collection',
```

```
    OPERATION    => NULL,
```

```
    WINDOW_NAME => 'FRIDAY_WINDOW');
```

```
END;
```

```
/
```



```
/* End of Script */
```

Verify the setup using the following SQL

```
SELECT WINDOW_NAME, OPTIMIZER_STATS  
FROM DBA_AUTOTASK_WINDOW_CLIENTS;
```

Should return

```
MONDAY_WINDOW    DISABLED  
TUESDAY_WINDOW   DISABLED  
WEDNESDAY_WINDOW DISABLED  
THURSDAY_WINDOW  DISABLED  
FRIDAY_WINDOW    DISABLED  
SATURDAY_WINDOW  ENABLED  
SUNDAY_WINDOW    ENABLED
```

5.2 Customizing Statistics Gathering for FLEXCUBE

The default statistics gathering is designed to be generic. It is recommended to customize the default statistics gathering to suit FLEXCUBE online and batch.

Following are the areas that would need customization for FLEXCUBE:

- [Statistics Histograms](#)

5.2.1 Statistics Histograms

Note the following:

- The default statistics gathering routine decides to collect histograms on specific tables based on certain criteria that are not documented.
- Statistics Histograms are not recommended for FLEXCUBE tables.

Configure the default statistics gathered without Histograms.

```
/* Start of Script – Script to be executed as SYS*/  
  
BEGIN  
  
    DBMS_STATS.SET_PARAM ('METHOD_OPT', 'FOR ALL COLUMNS SIZE 1');  
  
END;  
  
/  
  
/*End of Script */
```

Verify the setup using

```
SELECT DBMS_STATS.GET_PARAM ('METHOD_OPT') FROM DUAL;
```

Should return

```
FOR ALL COLUMNS SIZE 1
```

6. FLEXCUBE Database Storage Recommendations

Oracle database 10g release 2 onwards, Automatic Storage Management (ASM) is the recommended storage option for FLEXCUBE database. ASM is an integrated cluster aware volume manager and a file system designed and optimized for managing Oracle database files. ASM is the realization of the Oracle Stripe and Mirror Everything (SAME) storage management methodology researched and established as best practices for Oracle database environment over many years.

Note: For configuring ASM refer Automatic storage management best practice document provided by Oracle for your database version.

6.1 Key benefits of ASM

- I/O is spread evenly across all available disk drives to prevent hot spots and maximize performance.
- ASM eliminates the need for over provisioning and maximizes storage resource utilization facilitating database consolidation.
- Inherent large file support.
- Performs automatic online redistribution after the incremental addition or removal of storage capacity.
- Maintains redundant copies of data to provide high availability, or leverage 3rd party RAID functionality.
- Supports Oracle Database 19c as well as Oracle Real Application Clusters (RAC).
- Capable of leveraging 3rd party multipathing technologies.
- For simplicity and easier migration to ASM, an Oracle Database 19c database can contain ASM and non-ASM files. Any new files can be created as ASM files whilst existing files can also be migrated to ASM.
- RMAN commands enable non-ASM managed files to be relocated to an ASM disk group.
- Oracle Database 19c Enterprise Manager can be used to manage ASM disk and file management activities.

7. FLEXCUBE Database Backup Recommendations

Backup Policy is a very important ingredient of any High Availability system. Oracle recommends RMAN utility for database backup.

RMAN is acronym for Recovery Manager, is Oracle utility which will backup, restore, and recover oracle data files. RMAN is an Oracle provided utility for efficiently performing Backup and Recovery. RMAN is available as a part of the standard Installation and no separate installation is required.

Recovery Manager is a client/server application that uses database server sessions to perform backup and recovery. It stores metadata about its operations in the control file of the target database and, optionally, in a recovery catalog schema in an Oracle database.

You can invoke RMAN as a command-line executable from the operating system prompt or use some RMAN features through the Enterprise Manager GUI.

7.1 RMAN Vs Conventional Backup

- During a conventional hot backup, the amount of Redo generated during the backup would be more due to the fact that the redo logs during the hot backup store the entire block images rather than the change vectors.
- RMAN doesn't place the tablespace in a backup mode and hence the amount of Redo generated during the RMAN backup is considerably low.
- RMAN can identify block corruption during backup operations and RMAN supports Block recovery.
- RMAN automatically detects new data files and will backup them. Also, RMAN supports incremental backup method.
- RMAN backs up only the blocks that have been used at least once. Unused blocks are never backed up. Unused block here refers to the blocks where in the block header is zeroed
- RMAN enables us to test the backup without actually restoring the backup.
- RMAN can verify physical and logical structures of the database without actually performing backup.
- Usage of Shared Pool and Large Pool for RMAN
- RMAN uses DBMS_RCVMAN and DBMS_BACKUP_RESTORE packages for backup and recovery. These packages would be loaded in the shared pool for backup and restore operation. RMAN uses the PGA for backup and restore operation.
- RMAN Requires LARGE_POOL only if TAPE_IO_SLAVES and DBWR_IO_SLAVES are defined.
- **Sizing Large Pool** - $LARGE_POOL = (\text{Number of Channels}) * (16 \text{ MB} + \text{Tape Buffer})$

7.2 **Benefits of Using RMAN**

- RMAN is an intelligent tool that comes at no extra cost. It is available free with the Oracle Database.
- RMAN introduced in Oracle 8 it has become simpler with newer versions and easier than user managed backups.
- Provides proper security for Backups.
- You can be 100% sure your database has been backed up.
- Control file and Spfile of the database can be configured to be automatically backed up by RMAN.
- It contains detail of the backups taken etc. in its central repository Facility for testing validity of backups also commands like crosscheck to check the status of backup.
- Faster backups and restores compared to backups without RMAN.
- RMAN is the only native backup tool which supports incremental backups.
- Oracle 19c has got further optimized incremental backup which has resulted in improvement of performance during backup and recovery time.
- Parallel operations (Multiple Channels for Backup and Restore) are supported.
- Better querying facility for knowing different details of backup.
- No extra redo is generated when backup is performed, compared to conventional online backup.
- Maintains repository of backup metadata.
- Remembers backup set location.
- Knows what need to backed up.
- Knows what is required for recovery.
- Knows what backup are redundant.
- RMAN can back up the Database to Disk or directly to Tape. It is recommended that RMAN backup is performed to disk and then copied to tape.


7.3 Backup Strategy Recommendation

RMAN will not backup the below files so it is advised to take the copy of the below files on regular basis (weekly/any change/addition to the file).

- Tnsnames.ora
- Listener.ora
- Password file
- Init.ora

The Best practice is to take create the pfile once the spfile is updated.

Below is the recommended strategy.



8. Known Issues

After upgrading to Oracle database 19.6 from 19.3 base, pluggable database upon opening moves to restricted mode. This is a known issue. Oracle Support has released patch 30881407 for the same, which needs to be applied on top of 19.6.

9. Appendix

9.1 Script to Check Histograms on FLEXCUBE Schema

Following script would have to be executed in the FLEXCUBE schema:

```
select distinct table_name
from
(
select table_name from user_tab_columns where histogram!='NONE'
)
```

Should return **No Records**

9.2 Script to Remove Histograms on FLEXCUBE Schema

Following script would have to be executed in the FLEXCUBE schema if there are any rows:

```
declare
cursor cur_tables is
select distinct table_name
from
(
select table_name from user_tab_columns where histogram!='NONE'
);
begin
for rec_tables in cur_tables
loop
dbms_stats.gather_table_stats(ownname=>USER,tabname=>rec_tables.t
able_name,METHOD_OPT=>'FOR ALL COLUMNS SIZE
1',CASCADE=>TRUE,DEGREE=>2,ESTIMATE_PERCENT=>NULL);
end loop;
end;
```




FLEXCUBE Products Database Practices db 19C
[August] [2020]
Release 14.4.0.0.0

Oracle Financial Services Software Limited
Oracle Park
Off Western Express Highway
Goregaon (East)
Mumbai, Maharashtra 400 063
India

Worldwide Inquiries:
Phone: +91 22 6718 3000
Fax: +91 22 6718 3001
www.oracle.com/financialservices/

Copyright © [2013], [2020], Oracle and/or its affiliates. All rights reserved.

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate failsafe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

This software or hardware and documentation may provide access to or information on content, products and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.