

Oracle Financial Services Capital Adequacy Application Pack Performance Tuning Guide

Release 8.1.2.4.1

December 2023





Oracle Financial Services Basel Regulatory Compliance Analytics User Guide

Copyright © 2025 Oracle and/or its affiliates. All rights reserved.

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.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable.

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 fail-safe, 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.

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.

This software or hardware and documentation may provide access to or information about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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, except as set forth in an applicable agreement between you and Oracle.

For information on third party licenses, see the [OFSAALicensing Information User Manual](#).

Document Control

Version Number	Revision Date	Change Log
1.0	December 2023	This version is created to document the performance tuning steps in OFS CAAP.

Table of Contents

- 1 Preface5**
 - 1.1 Scope of the Guide 5
 - 1.2 Intended Audience 5
 - 1.3 Access to Oracle Support 5
 - 1.4 Related Information Sources 5
- 2 Introduction6**
- 3 OFSAAI Application Performance Tuning.....7**
- 4 Database Performance Tuning.....9**
 - 4.1 Mandatory step 9
 - 4.2 Optional Steps 10
- 5 Operating System Level Tuning Parameters15**
 - 5.1 Optional Steps 15
- OFSAA Support.....17**
- Send Us Your Comments 18**

1 Preface

This chapter provides a brief description of the scope, the audience, and the references for this document.

1.1 Scope of the Guide

This document includes the performance tuning steps for OFS Capital Adequacy Application Pack.

1.2 Intended Audience

This manual is intended for database administrators and application administrators.

1.3 Access to Oracle Support

Oracle customers have access to electronic support through [My Oracle Support](https://support.oracle.com/signin). For information, visit support.oracle.com/signin.

Or, visit [Accessibility Learning and Support | Oracle](https://accessibility.oracle.com) if you are hearing impaired.

1.4 Related Information Sources

We strive to keep this and all other related documents updated regularly. Visit the below links to download the latest version available.

- [OHC Documentation Library](#) for OFS CAP Application Pack
- [OHC Documentation Library](#) for OFS AAAI Application Pack

2 Introduction

This document includes the performance tuning steps for OFS Capital Adequacy Application Pack (CAAP). Follow the steps in this document to achieve the batch performance across all jurisdictions run definitions.

 Note

The performance tuning steps mentioned in this document were implemented and tested during the development and Quality Assurance (QA) phase of the OFS CAP Release 8.1 development life cycle.

3 OFSAAI Application Performance Tuning

This section includes the performance tuning steps for OFSAAI. Perform the following steps in the OFSAAI instance:

1. Navigate to the \$FIC_HOME/ficdb/conf directory and modify the OracleDB configuration (OracleDB.conf) file as follows:
CNF_PARALLEL_QUERY=ENABLE
CNF_PARALLEL_DML=ENABLE
CNF_DEGREE_OF_PARALLELISM=
CNF_PARALLEL_DEGREE_POLICY=MANUAL

Decide the value for CNF_DEGREE_OF_PARALLELISM based on the number of CPU cores available in the database (DB) engine server. It is recommended to use 40 % of the available core.

For example, if the CPU cores are 10 in count, then assign CNF_DEGREE_OF_PARALLELISM=4.

Note

The setting of these parameters assumes that the database is enabled for Parallel DML execution

2. Modify the .profile file of the OFSAAI instance user.
 Increase the memory setting for the variable X_ARGS_RLEXE as follows:
 X_ARGS_RLEXE="-Xms512m -Xmx4g -XX:+UseAdaptiveSizePolicy -XX:MaxPermSize=1024M -XX:+UseParallelOldGC -XX:+DisableExplicitGC
\$X_ARGS_GEN"
 The Rule engine uses these settings. If the OFSAAI application server has a memory of 32 GB, this value can be set up to 8 GB. For example: -Xms4g -Xmx8g
3. In the ETLLoader.sh file, used by the T2T, increase the load of Xms6400m -Xmx10256m to 10G.
4. Modify the SETUP_MASTER table as follows:
 - a. Log in to the OFSAAI Atomic schema (this is the schema where the CAP objects are physicalized).
 - b. Verify the following Component Codes and Component Values in the SETUP_MASTER table.

Table 1: Component Codes and Component Values in the SETUP_MASTER Table

V_COMPONENT_CODE	V_COMPONENT_DESC	V_COMPONENT_VALUE
DT_PARALLEL_DOP	Degree of parallelism to be used in DML and Queries statements in the data transformations.	4
DT_PARALLEL_ENABLE	Enables parallel sessions for DML and Queries statements in data transformations.	Y

 Note

- Decide the Parallel 4 value based on the number of CPU cores available in the database engine server. It is recommended to use 40 % of the available core. For example, if the CPU cores are 10 in count, then assign parallel to 4.
- The setting of these parameters assumes that the database is enabled for Parallel DML execution.

5. Modify the CONFIGURATION table as follows:
- a. Login to the Config schema.
 - b. Set the PARAMVALUE settings for the PARAMNAME as follows.

Table 2: PARAMNAME and PARAMVALUE

PARAMNAME	PARAMVALUE	DESCRIPTION
QRY_OPT_USE_ROWID	Y	To use ROWID in MERGE ON statement for RRF, specify Y otherwise N.

6. Increase the Connection Pool in the WebLogic console and set the Connection Pool to 100.

4 Database Performance Tuning

This section includes the database performance tuning steps.

4.1 Mandatory step

Perform the following steps to create table partitions using the provided scripts.

1. Partition tables:

- FSI_MULTIPLE_RATING_PROCESSING: Partition by list on the V_RECORD_TYPE column:

```
ALTER TABLE FSI_MULTIPLE_RATING_PROCESSING MODIFY
PARTITION BY LIST ("V_RECORD_TYPE")
(PARTITION BNK_NON_SEC_EXP VALUES ('BNK_NON_SEC_EXP'),
PARTITION BNK_NON_SEC_ULTY VALUES ('BNK_NON_SEC_ULTY'),
PARTITION BNK_SEC_ULTY VALUES ('BNK_SEC_ULTY'),
PARTITION DRV_NON_SEC_ULTY VALUES ('DRV_NON_SEC_ULTY'),
PARTITION DRV_SEC_ULTY VALUES ('DRV_SEC_ULTY'),
PARTITION INV_MR_EXP VALUES ('INV_MR_EXP'),
PARTITION INV_MR_ULTY VALUES ('INV_MR_ULTY'),
PARTITION INV_NON_SEC_ULTY VALUES ('INV_NON_SEC_ULTY','INV_NON_SEC_EXP'),
PARTITION INV_SEC_ULTY VALUES ('INV_SEC_ULTY'),
PARTITION OTH VALUES ('OTH'),
PARTITION SFT_NON_SEC_EXP VALUES ('SFT_NON_SEC_EXP'),
PARTITION SFT_NON_SEC_ULTY VALUES ('SFT_NON_SEC_ULTY'),
PARTITION SFT_SEC_ULTY VALUES ('SFT_SEC_ULTY'),
PARTITION SFT_PLACED_COLL_EXP VALUES ('SFT_PLACED_COLL_EXP', 'SFT_MITIGANT'),
PARTITION MITIGANT VALUES ('MITIGANT'),
PARTITION DRV_NON_SEC_EXP VALUES ('DRV_NON_SEC_EXP'),
PARTITION OTH_PLACED_COLL_EXP VALUES ('OTH_PLACED_COLL_EXP'));
```

- FSI_CAP_BANKING_EXPOSURES: Partition by list on the V_RECORD_TYPE column:

```
ALTER TABLE FSI_CAP_BANKING_EXPOSURES MODIFY
PARTITION BY LIST ("V_RECORD_TYPE")
(PARTITION BNK_NON_SEC_EXP VALUES ('BNK_NON_SEC_EXP'),
PARTITION BNK_NON_SEC_ULTY VALUES ('BNK_NON_SEC_ULTY'),
PARTITION BNK_SEC_ULTY VALUES ('BNK_SEC_ULTY'));
```

4.2 Optional Steps

Perform the following steps after completing the mandatory step.

1. Create the following indexes:

Note

If the creation of the following indexes fails due to any reason – such as *object name already exists*, then provide new index names.

Scripts

```
create index IDX$$_024B0001 on FCT_ENTITY_INFO("N_BASEL_CAP_CONSL_APPR_SKEY");
create index IDX$$_024B0002 on FSI_CAP_PARTY_FINANCIALS("N_PARTY_SKEY");
create index IDX$$_024B0003 on DIM_LIEN_POSITION ("D_RECORD_START_DATE","V_LIEN_POSITION_CODE", "D_RECORD_END_DATE");
create index IDX$$_024B0004 on FSI_PROD_TXN_TYPE_RECLASS("N_BASEL_PROD_TYPE_SKEY");
create index IDX$$_024B0005 on
DIM_INSTRUMENT_CONTRACT("D_RECORD_START_DATE","V_INSTRUMENT_CODE","D_RECORD_END_DATE");
create index IDX$$_02580001 on FSI_RUN_PARAMETER("N_RUN_SKEY");
create index IDX$$_02580002 on RUN_PARAMETERS("N_MIS_DATE_SKEY");
create index IDX$$_02580003 on STG_SWAPS_CONTRACTS("V_PRODUCT_CODE");
create index IDX$$_02580004 on STG_SWAPS_CONTRACTS("V_LV_CODE");
create index IDX$$_02580005 on FCT_ENTITY_INFO("V_ENTITY_CODE");
create index IDX$$_02580006 on DIM_NETTING_AGREEMENT("V_NETTING_AGREEMENT_CODE");
create index IDX$$_02580007 on DIM_UNDERLYING("V_UNDERLYING_CODE");
create index IDX$$_02580008 on DIM_INSTRUMENT_CONTRACT("V_INSTRUMENT_CODE");
create index IDX$$_02580009 on DIM_CLEARED_TXN_BANK_ROLE("V_CLEARED_TXN_BANK_ROLE_CODE");
create index IDX$$_02590001 on STG_LOAN_CONTRACTS("V_PROD_CODE");
create index IDX$$_02590002 on FSI_CAP_CURRENCY_CONVERSION("V_CCY_CODE");
create index IDX$$_02590003 on DIM_INTEREST_TYPE("V_INTEREST_TYPE");
CREATE INDEX IDX$$_02590004 ON STG_CARDS ("V_PROD_CODE");
create index IDX$$_025F0001 on STG_UNDERLYING_EXPOSURES("V_PARENT_EXPOSURE_ID","V_GAAP_CODE");
create index IDX$$_025F0002 on FSI_CAP_ACCOUNT_PORTFOLIO_MAP("N_CAP_APPROACH_TYPE_SKEY");
create index IDX$$_025F0003 on DIM_HOLDING_TYPE("V_HOLDING_TYPE");
create index IDX$$_025F0004 on FSI_SENIORITY_RECLASSIFICATION("N_RUN_SKEY");
create index IDX$$_025F0005 on FSI_PARTY_TYPE_CLASSIFICATION("N_RUN_SKEY");
create index IDX$$_03630001 on FSI_PARTY_TYPE_CLASSIFICATION("N_MIS_DATE_SKEY");
create index IDX$$_025F0006 on FSI_PARTY_TYPE_SECTOR_RECLASS("N_RUN_SKEY");
create index IDX$$_025F0007 on FSI_CAP_INDUS_SECTOR_RECLASS("N_RUN_SKEY");
create index IDX$$_02600001 on FCT_ENTITY_INFO("N_RUN_SKEY","F_CAP_CONSL_ENTITY_IND");
create index IDX$$_02600002 on DIM_PARTY_ROLE("V_PARTY_ROLE_CODE");
create index IDX$$_02610001 on STG_OPTION_CONTRACTS("V_PRODUCT_CODE");
create index IDX$$_02610002 on STG_OPTION_CONTRACTS("V_LV_CODE");
create index IDX$$_026B0001 on STG_UNDERLYING_EXPOSURES("V_PROD_CODE");
create index IDX$$_026B0002 on STG_UNDERLYING_EXPOSURES("V_LV_CODE");
create index IDX$$_026B0003 on FCT_INSTRUMENT_CONTRACT_DTL("N_INSTRUMENT_SKEY");
create index IDX$$_026B0004 on DIM_BANK_INSTRUMENT_TYPE("V_BANK_INSTR_TYPE_CODE");
create index IDX$$_026B0006 on FSI_CAP_PARTY_FINANCIALS("N_MIS_DATE_SKEY");
create index IDX$$_02840001 on SEC_POOL_MASTER("V_SEC_POOL_ID","D_RECORD_START_DATE");
```

```
create index IDX$$_028E0001 on FSI_SENIORITY_RECLASSIFICATION("N_MIS_DATE_SKEY");
create index IDX$$_028F0001 on FSI_CAP_INDICATOR("V_CAP_INDICATOR_TYPE");
create index IDX$$_028F0002 on FSI_CAP_BANDS("N_CAP_BAND_LOWER_BOUND_VALUE","N_CAP_BAND_UPPER_BOUND_VALUE");
create index IDX$$_05500002 on FSI_BASEL_ASSET_CATEGRY_MAP("N_RUN_SKEY","N_MIS_DATE_SKEY");
create index IDX$$_05500007 on FSI_BASEL_RATING_BAND_MAP("N_RUN_SKEY");
```

2. Alter the following tables for NOLOGGING:

```
ALTER TABLE FSI_CAP_BANKING_EXPOSURES NOLOGGING;
ALTER TABLE FSI_CAP_INVESTMENT_EXPOSURES NOLOGGING;
ALTER TABLE FSI_CAP_SFT_EXPOSURES NOLOGGING;
ALTER TABLE FSI_CAP_DERIVATIVES NOLOGGING;
ALTER TABLE FSI_CAP_MITIGANTS NOLOGGING;
```

3. Execute the following updates statements for T2Ts to pick hints in select and insert clauses:

```
UPDATE AAI_DMT_MAPPING_INFO SET V_HINT = ' APPEND ';
UPDATE AAI_DMT_DEFINITION SET V_HINT = ' PARALLEL (4) ';
UPDATE aai_dmt_mapping_info set V_PRESCRIPT = 'ALTER SESSION SET STAR_TRANSFORMATION_ENABLED=TRUE; ALTER SESSION
FORCE PARALLEL DML PARALLEL 4; ALTER SESSION FORCE PARALLEL QUERY PARALLEL 4; ALTER SESSION SET
OPTIMIZER_ADAPTIVE_FEATURES = FALSE;';
UPDATE aai_dmt_definition set V_PRESCRIPT = 'ALTER SESSION SET STAR_TRANSFORMATION_ENABLED=TRUE; ALTER SESSION FORCE
PARALLEL DML PARALLEL 4; ALTER SESSION FORCE PARALLEL QUERY PARALLEL 4; ALTER SESSION SET
OPTIMIZER_ADAPTIVE_FEATURES = FALSE;';
```

 Note

- Decide the Parallel 4 value based on the number of CPU cores available in the database engine server. It is recommended to use 40 % of the available core. For example, if the CPU cores are 10 in count, then assign parallel to 4.

4. Execute the following query to append the hints for rules:

```
INSERT INTO AAI_OBJ_QUERY_OPTIMIZATION(V_OBJ_CODE, V_INFODOM_CODE, V_OBJ_TYPE, V_EXECUTION_MODE, F_USE_PARTITION,
F_USE_ROWID, V_MERGE_HINT, V_SELECT_HINT, V_PRE_SCRIPT, V_POST_SCRIPT) SELECT V_RULE_NAME, V_INFODOM_NAME, 'RL',
'MERGE', 'Y', F_USE_ROWID, NULL, 'PARALLEL(4)', 'ALTER SESSION SET STAR_TRANSFORMATION_ENABLED=TRUE; ALTER SESSION
FORCE PARALLEL DML PARALLEL 4; ALTER SESSION FORCE PARALLEL QUERY PARALLEL 4; ALTER SESSION SET
OPTIMIZER_ADAPTIVE_FEATURES = FALSE;', NULL FROM PR2_RULE_B;
```

 Note

- Decide the Parallel 4 value based on the number of CPU cores available in the database engine server. It is recommended to use 40 % of the available core. For example, if the CPU cores are 10 in count, then assign parallel to 4.

5. Execute the result of below query to disable FK constraints on all FSI and FCT

```
select 'alter table ' || TABLE_NAME || ' disable constraint ' || CONSTRAINT_NAME || ';' from user_constraints where CONSTRAINT_TYPE = 'R' and
(TABLE_NAME like 'FSI%' or TABLE_NAME like 'FCT%');
```

6. Gather Stats for all the tables in the Atomic Schema periodically.

You must gather statistics periodically for objects where the statistics become stale over time, due to changing data volumes or changes in column values. New statistics should be gathered after a schema object's data or structure is modified in ways that make the previous statistics inaccurate. For example, after loading a significant number of rows into a table, collect new statistics on the number of rows. After updating data in a table, collecting new statistics on the number of rows is not required, but you might need new statistics on the average row length.

Note

Modify the ownname **<Atomic_Schema>** to the actual schema name used by the Atomic schema.

Use the DBMS_STATS package to generate the statistics. Following are examples of Gather Stats:

```
execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'RUN_PARAMETERS', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);

execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_MULTIPLE_RATING_PROCESSING',
estimate_percent => DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);

execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_CAP_BANKING_EXPOSURES', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);

execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_CAP_CURRENCY_CONVERSION', estimate_percent
=> DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);

execute dbms_stats.gather_index_stats(ownname => '<ATOMIC_SCHEMA>', indname => 'IDX$$_015C0013', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE);

execute dbms_stats.gather_index_stats(ownname => '<ATOMIC_SCHEMA>', indname => 'IDX$$_015C000C', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE);

execute dbms_stats.gather_index_stats(ownname => '<ATOMIC_SCHEMA>', indname => 'FCT_ENTITY_INFO_IDX1', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE);

execute dbms_stats.gather_index_stats(ownname => '<ATOMIC_SCHEMA>', indname => 'FCT_ENTITY_INFO_IDX2', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE);

execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_RW_MAP_MASTER', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);

execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_ISSUE_RATING_DETAILS', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);

execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'STG_ACCOUNT_MITIGANT_MAP', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);

execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'STG_LOAN_CONTRACTS', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);

execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'STG_UNDERLYING_EXPOSURES', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);

execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'STG_INVESTMENTS', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);

execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'STG_LC_CONTRACTS', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);

execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'STG_SWAPS_CONTRACTS', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);

execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'STG_FX_CONTRACTS', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);
```

```
execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'STG_MM_CONTRACTS', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);
execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_INSTR_TYPE_CLASSIFICATION', estimate_percent
=> DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);
execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_INSTR_TYPE_ULTY_TYPE_RECLS', estimate_percent
=> DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);
execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_CAP_RUN_EXE_PARAMETERS', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);
execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_PROD_TXN_TYPE_RECLASS', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);
execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_SENIORITY_RECLASSIFICATION', estimate_percent
=> DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);
execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_PARTY_TYPE_CLASSIFICATION', estimate_percent
=> DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);
execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_REG_PROD_TYPE_RECLASS', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);
execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_CAP_PARTY_FINANCIALS', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);
execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FCT_ENTITY_INFO', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);
execute dbms_stats.gather_table_stats(ownname => '<ATOMIC_SCHEMA>', tabname => 'FSI_CAP_ACCOUNT_LOAD_RUN_MAP', estimate_percent =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL COLUMNS SIZE AUTO', cascade => TRUE);
```

7. It is recommended, that before starting the batches, **collect system level stats as well.**
Below are the sample expected values for Exadata environment.

ulimit -a		
• core file size (blocks)	(-c)	0
• data seg size (kbytes)	(-d)	unlimited
• scheduling priority	(-e)	0
• file size (blocks)	(-f)	unlimited
• pending signals	(-i)	2900764
• max locked memory (kbytes)	(-l)	32768
• max memory size (kbytes)	(-m)	unlimited
• open files	(-n)	1024
• pipe size (512 bytes)	(-p)	8
• POSIX message queues (bytes)	(-q)	819200
• real-time priority	(-r)	0
• stack size (kbytes)	(-s)	10240
• cpu time (seconds)	(-t)	unlimited
• max user processes	(-u)	4096
• virtual memory (kbytes)	(-v)	unlimited
• file locks	(-x)	unlimited

8. Session level parameters:

- Database parameters can be tuned to get the best optimal results from the optimizer, which help change the relative costs of a full table scan versus index operations.
- Memory parameters can be updated for faster output. These help clearing latches and library cache, better read and write operations, along with other events like clear direct path write and read etc.
- Caching parameters can be updated along with related parameters to maintain a pool for buffers of the primary block size. These also help in events like clear slave session stats and buffer busy waits.

Parameters	Values
optimizer_index_caching	0
optimizer_index_cost_adj	100
sga_max_size	300G
sga_target	300G
pga_aggregate_target	160G
pga_aggregate_limit	0
session_cached_cursors	74000
cursor_sharing	FORCE
open_cursors	64000
shared_pool_size	80G
cpu_count	100
processes	6000
sessions	7600
parallel_force_local	FALSE
db_recovery_file_dest_size	64000G
undo_management	AUTO
gcs_server_processes	10

5 Operating System Level Tuning Parameters

This section includes the Operating System Level Tuning Parameters steps.

5.1 Optional Steps

Perform the following steps:

Below parameters needs to be set on the application server:

These parameters help in improving soft/hard parsing and help overcome any application binary specific issues while parsing the code/data at the OS level.

These includes system parameters which decide

- how many files can be kept open by a process
- how large file can be created by a user
- how much memory can be used by the different components of the process such as the stack, data and text segments.

- Add below entries in "/etc/security/limits.conf"

*	soft	memlock	268435456
*	hard	memlock	268435456
*	soft	msgqueue	268435456
*	hard	msgqueue	268435456
*	hard	stack	1048576
*	soft	stack	1048576
oraofss	soft	nofile	65536
oraofss	hard	nofile	65536
oraofss	soft	nproc	15000

- ulimit -s
1048576 (or ulimit –s unlimited)
- Set system limits on application server system, for the application’s OS level user, so that they match the following values

ulimit -a

- core file size (blocks) (-c) 0
- data seg size (kibytes) (-d) unlimited
- scheduling priority (-e) 0
- file size (blocks) (-f) unlimited
- pending signals (-i) 2061849
- max locked memory (Kibytes) (-l) 64
- max memory size (Kibytes) (-m) unlimited
- open files (-n) 1024
- pipe size (512 bytes) (-p) 8

- POSIX message queues (bytes) (-q) 819200
- real-time priority (-r) 0
- stack size (kbytes) (-s) 8192
- cpu time (seconds) (-t) unlimited
- max user processes (-u) 2061849
- virtual memory (kbytes) (-v) unlimited
- file locks (-x) unlimited

OFSAA Support

Raise a Service Request (SR) in [My Oracle Support \(MOS\)](#) for queries related to the OFSAA applications.

Send Us Your Comments

Oracle welcomes your comments and suggestions on the quality and usefulness of this publication. Your input is an important part of the information used for revision.

- Did you find any errors?
- Is the information clearly presented?
- Do you need more information? If so, where?
- Are the examples correct? Do you need more examples?
- What features did you like most about this manual?

If you find any errors or have any other suggestions for improvement, indicate the title and part number of the documentation along with the chapter/section/page number (if available) and contact the My Oracle Support.

- Before sending us your comments, you might like to ensure that you have the latest version of the document wherein any of your concerns have already been addressed. You can access My Oracle Support site that has all the revised or recently released document.