

# Oracle 11G New Features

Major contributors:  [Deepak Mishra](#) | [All Contributors](#)

## Sommaire

- SQL Tuning Improvements in Oracle 11G ..... 2
  - These are some of the enhanced Oracle 11g SQL tuning capabilities: ..... 2
  - Result Caches..... 3
  - New System/Performance Testing Tools in Oracle 11G ..... 3
    - Database Replay ..... 3
  - SQL Performance Analyzer (SPA) ..... 4
- Repair Advisors..... 4
  - Fault Diagnostics in Oracle 11G..... 4
    - Automatic Diagnostic Repository ..... 4
    - Automatic Health Monitoring ..... 4
    - Hang Manager ..... 4
    - Incident Packaging Service ..... 4
    - Support Workbench ..... 5
- Flashback Enhancements in Oracle 11G ..... 5
  - Flashback Data Archive ..... 5
  - Flashback Transaction ..... 5
- Automatic Memory Tuning ..... 5
- Some of the Administration related features are as follows: ..... 6
  - Virtual columns..... 6
  - Faster sorting..... 6
  - Data Compression ..... 6
  - New partitioning capabilities in Oracle 11G..... 7

Partitioning Matrix .....	7
Enhancements in Database Security .....	7
Case Sensitivity .....	7
Strong Authentication .....	7
Backwards Compatibility .....	7
Administration.....	8
Disabling Password Case Sensitivity .....	8
Password File.....	8

The new Oracle 11g database has more than 482 new features. Not all new features have been fully documented yet; I have mentioned some of the enhancements made to oracle 10G database.

## SQL Tuning Improvements in Oracle 11G

---

As an Oracle DBA you would have experienced the advice for SQL performance improvements that Oracle Database 10g's SQL Tuning Advisor and SQL Access Advisor provide.

### These are some of the enhanced Oracle 11g SQL tuning capabilities:

---

SQL Access Advisor can now make recommendations on how partitioning might be applied to existing tables, indexes, and materialized views to improve an application's performance. SQL statements can now tune themselves automatically, this has come up as a revolution to an oracle databases.

The 10g automatic tuning advisor makes tuning suggestions in the form of SQL profiles that will improve performance.

You can tell 11g to automatically apply SQL profiles for statements where the suggested profile gives 3-times better performance than the existing statement.

The performance comparisons are done by a new administrative task during a user-specified maintenance window.

Statistics for the Cost-Based Optimizer (CBO) are now published separately from being gathered.

- This means that recomputed statistics for the CBO will not necessarily cause existing cursors to become invalidated.

Oracle Database 11g now supports retention of historical execution plans for a SQL statement.

- This means that the CBO can compare a new execution plan against the original plan and, if the old plan still offers better performance than the new one, it can decide to continue to use the original execution plan.

Statistics can be gathered for multiple columns. Collecting statistics for multiple columns gives the CBO the ability to more accurately select rows based on common multi-column conditions or joins.

## Result Caches

---

As a DBA / Application Developer I have thought that Oracle database would provide a method to retain the result set in memory from a complex query. These are constant data that hardly ever change, but must still be read and used across multiple applications - for example, a list of all Zip Codes and their corresponding Area names for lookup when processing addresses for new customers.

Oracle Database 11g fills this gap with three new structures called result caches, and each structure has a different purpose:

The SQL query result cache is an area of memory in the Shared Global Area (SGA) that can retain the result sets that a query generates.

The PL/SQL function result cache can store the results from a PL/SQL function call.

Finally, the OCI client result cache can retain results from queries or functions on the application server from which the OCI call originated.

By retaining result sets in these in-memory caches, the results are immediately available for reuse by any user session.

For user sessions that connect to the database through an application server, the OCI client cache permits those sessions to simply share the results that are already cached on the application server without having to reissue a query directly against the database.

These result caches eliminate unnecessary “round trips” to the database server to collect relatively static reference data. This brings a great improvement in overall database throughput.

For Example:

The `/*+result cache*/` SQL hint - This suggests that the result data will be cached in the data buffers, and not the intermediate data blocks that were accessed to obtain the query results.

You can cache SQL and PL/SQL results for super-fast subsequent retrieval.

```
SQL> Alter table table_name cache results;
```

## New System/Performance Testing Tools in Oracle 11G

---

As a DBA, the most difficult task which I have found is diagnosing the database performance issues and tuning them.

Also the patching the database's application code, database patch set, or hardware configuration will affect that database's performance, which usually meant purchasing a third-party package (e.g. Mercury Interactive's LoadRunner) to generate a sample workload against the database using the next version of the application code, and then comparing the results against baseline performance for the current application code version.

This can be very useful for oracle regression testing.

So there is good news for Tuning Engineer that, Oracle Database 11g has come to the rescue these issues with two new utilities that offer a great advantage which performing the system/performance testing:

### Database Replay

Database Replay can capture generated workloads from production systems at the database level. Therefore, it's no longer necessary to run actual application code to duplicate the load on the database, and this also improves accuracy of the simulated workload because it limits or removes

other factors like network latency.

These captured workloads can then be replayed on a quality assurance i.e. UAT database so that the impact of application changes, software patches, and even hardware upgrades can be measured accurately. This feature is especially valuable in detecting performance issues that could potentially cause a production database performance that might go otherwise undetected until the changes have been deployed.

## SQL Performance Analyzer (SPA)

---

The SQL Performance Analyzer enhances the existing Oracle Database 10g SQL tuning components.

The SPA provides the ability to capture a specific SQL workload in a SQL Tuning Set, take a performance baseline before a major database or system change, make the desired change to the system, and then replay the SQL workload against the modified database or configuration.

The before and after performance of the SQL workload can then be compared with just a few clicks of the mouse. The DBA only needs to find any SQL statements that are now performing poorly and tune them via the SQL Tuning Advisor.

## Repair Advisors

---

Oracle 10G introduced some of the database performance advisory framework like Undo Advisor, Segment Advisor, SQL Tuning Advisor, SQL Access Advisor, MTTR Advisor, and the best of all Automatic Database Diagnostic Monitor (ADDM). Oracle Database 11g expands this advisory framework by introducing a new Database Repair Advisor.

The main goal of this new Advisor is to locate root causes of a failure, identify and present options for repairing these root causes, and correct the identified problems with self-healing mechanisms.

## Fault Diagnostics in Oracle 11G

---

Oracle Database 11g adds some of improved fault diagnostics mechanism to make it extremely easy for DBA's who have little exposure of system/database can easily detect and quickly resolve problems using this advance feature. Here are some of the points mentioned:

### Automatic Diagnostic Repository

The Automatic Diagnostic Repository (ADR) is at the heart of Oracle Database 11g's new fault diagnostic framework. The ADR is a central, file-based repository external to the database itself, and it's composed of the diagnostic data – alert logs, core dumps, background dumps, and user trace files - collected from individual database components from the first moment that a critical error is detected. When critical errors are detected, Oracle automatically creates an "incident" ticket, notifying the DBA instantly.

### Automatic Health Monitoring

When a problem within the database is detected, the new Health Monitor (HM) utility will automatically perform a series of integrity checks to determine if the problem can be traced to corruption within database blocks, redo log blocks, undo segments, or dictionary table blocks. HM can also be fired manually to perform checks against the database's health on a periodic basis.

### Hang Manager

Oracle Database 10g introduced the Hang Analysis tool in Enterprise Manager, and Oracle Database 11g now expands this concept with the Hang Manager.

### Incident Packaging Service

If the problem can't be solved using these tools, it may be time to ask for help from Oracle Support.

The new Incident Packaging Service (IPS) facility provides tools for gathering and packaging all necessary logs that Oracle Support typically needs to resolve a Service Request.

## Support Workbench

Though it's stored outside of the database itself, the ADR can be accessed via either Enterprise Manager or command-line utilities.

Once the ADR has detected and reported a critical problem, the DBA can interrogate the ADR; report on the source of the problem, and in some cases even implement repairs through the Support Workbench, a new facility which comes with Enterprise Manager.

# Flashback Enhancements in Oracle 11G

---

Oracle Database 10g database has come up flashback feature to recover the database with the ability to perform an incomplete recovery of the database with Flashback Database enhancement.

Oracle Database 10g also provided four new logical database recovery features: Flashback Table, Flashback Drop, Flashback Version Query, and Flashback Transaction Query.

Oracle Database 11g expands this new feature by introducing two new Flashback features:

## Flashback Data Archive

---

This new feature offers the ability to retain the reciprocal UNDO information for critical data significantly beyond the point in time that it would be flushed out of the UNDO tablespace.

Therefore, it's now possible to hold onto these reciprocal transactions essentially indefinitely.

Once this feature is enabled, all retained transaction history can be viewed, and this eliminates the task of creating corresponding history tracking tables for critical transactional tables.

And as you might expect, Oracle Database 11g also provides methods to automatically purge data retained in the Flashback Data Archive once a specified retention period has been exceeded.

## Flashback Transaction

---

This is an extension of the Flashback Transaction Query functionality introduced in Oracle Database 10g, Flashback Transaction allows the DBA to back out of the database one or more transactions.

As well as any corresponding dependent transactions by applying the appropriate UNDO statements for the affected transaction's to the corresponding affected rows in the database.

# Automatic Memory Tuning

---

In 11g, all memory can be tuned automatically by setting one parameter. Prior to Oracle 11g, the DBA set the `sga_target` and `sga_max_size` parameters, allowing Oracle to reallocate RAM within the SGA. The PGA was independent, as governed by the `pga_aggregate_target` parameter.

Now in Oracle 11g we see the `memory_max_target` parameter which governs the total maximum RAM for both the PGA and SGA regions and the new `memory_target` parameter which governs the existing sizes. This allows RAM to be de-allocated from the SGA and transferred to the PGA.

As of Oracle 10g, Automatic Memory Management (AMM) only allows shifting of RAM within the confines of `sga_max_size`. Now in Oracle 11g, we see this new Automatic Shared Memory Management (ASMM).

The `memory_target` parameter is dynamic (changeable with "alter system" commands), whereby RAM can be de-allocated from an instances SGA/PGA and re-allocated to another instance.

`MEMORY_TARGET` specifies the Oracle system-wide usable memory. The database tunes memory to the `MEMORY_TARGET` value, reducing or enlarging the SGA and PGA as needed. In a text initialization parameter file, if you omit the line for `MEMORY_MAX_TARGET` and include a value for `MEMORY_TARGET`, the database automatically sets `MEMORY_MAX_TARGET` to the value of

MEMORY\_TARGET.

If you omit the line for MEMORY\_TARGET and include a value for MEMORY\_MAX\_TARGET, the MEMORY\_TARGET parameter defaults to zero.

After startup, you can then dynamically change MEMORY\_TARGET to a non-zero value, provided that it does not exceed the value of MEMORY\_MAX\_TARGET.

## Some of the Administration related features are as follows:

### Virtual columns

---

Oracle 11g virtual table columns are columns that are actually functions.

For Example you can create table as follows:

```
SQL>create table t1 (c1 number, c2 number, c3 as (c1+c2) virtual);
```

### Faster sorting

---

Starting in 10gr2 we see an improved sort algorithm, "Oracle10gRw introduced a new sort algorithm which is using less memory and CPU resources. A hidden parameter \_newsort\_enabled = {TRUE|FALSE} governs whether the new sort algorithm will be used."

### Data Compression

---

Compression technology that can reduce the amount of storage required for data sets by up to two thirds.

Oracle uses a unique compression algorithm specifically designed to work with relational data. The algorithm works by eliminating duplicate values within a database block, even across multiple columns.

Compressed blocks contain a structure called a symbol table that maintains compression metadata. When a block is compressed, duplicate values are eliminated by first adding a single copy of the duplicate value to the symbol table.

Each duplicate value is then replaced by a short reference to the appropriate entry in the symbol table. Through this innovative design, compressed data is self-contained within the database block as the metadata used to translate compressed data into its original state is contained within the block.

When compared with competing compression algorithms that maintain a global database symbol table, Oracle's unique approach offers significant performance benefits by not introducing additional I/O when accessing compressed data.

One significant advantage is Oracle's ability to read compressed blocks directly without having to first uncompress the block. Therefore, there is no measurable performance degradation for accessing compressed data.

In fact, in many cases performance may improve due to the reduction in I/O since Oracle will have to access fewer blocks. Further, you can achieve a more efficient buffer cache by storing more data in the cache without having to increase your memory.

Rman Backups and Datapump exports can also be compressed. Datapump exports can be compressed using the following four options available with compress clause:

1. All.
2. Data-only.

3. Metadata-only.
4. None.

## New partitioning capabilities in Oracle 11G

---

Partitioning by parent/child references.

Partitioning by virtual columns.

More composite partitioning options, including

- o Range/Range.
- o List/Range.
- o List/Hash.
- o List/List.

Interval partitioning, which allows you to automatically create new partitions based on intervals, such as every month or every day.

In previous versions of Oracle, Composite Partitioning, or the use of subpartition, was permitted but the partition/subpartition combinations were limited.

In Oracle Database 11g, these restrictions have been removed to allow for all combinations of Composite Partitions, including the new Interval Partitions.

## Partitioning Matrix

---

	Range	List	Hash
Range	YES	YES	YES
List	YES	YES	YES
Interval	YES	YES	YES

## Enhancements in Database Security

---

Passwords can be Case Sensitive from Oracle 11G onwards.

### Case Sensitivity

11g will support case sensitive passwords automatically, when a user is created with the familiar syntax, Oracle will automatically store a new type of password verifier (hash value) in the data dictionary table that holds the user information. The same happens when a user or DBA changes an existing password.

### Strong Authentication

The verifier is based on 160 bit SHA-1 hashing algorithm, this means that in addition to allow for case sensitivity, the hashing algorithm was changed to create a longer and more secure type of verifier that is accepted generally by the industry and which is employed in several widely used security applications and protocols.

### Backwards Compatibility

When a database is migrated, only the hash values based on the legacy password hashing algorithm are available, therefore the passwords will remain case insensitive until they are changed.

It is recommended to change all database passwords soon after the migration to 11g as to take advantage of the added security that the case sensitive passwords and the new SHA-1 hashing provide.

A DBA should change the administrative accounts such as the SYS, SYSTEM etc. immediately and can use a password management policy (Profile) to force end-user password changes within a reasonable time.

## Administration

---

The view DBA\_USERS has new column PASSWORD\_VERSIONS rendered as follows:

decode(length(u.password),16,'10G ',NULL)||NVL2(u.spare4, '11G ',NULL) for example:

```
SQL> select USERNAME, PASSWORD_VERSIONS from dba_users;
```

```
USERNAME PASSWORD
```

---

```
-----
```

```
... SCOTT 10G 11G
```

In this case it means both old and new-style hash values are available for user SCOTT, the new hash value is stored in the USER\$.SPARE4 column, as long as this remains NULL it means the password has not been changed since the migration and the user will have the old case insensitive password.

## Disabling Password Case Sensitivity

---

When you create or modify user accounts, by default, passwords are case sensitive. To control the use of case sensitivity in passwords for application-to-password file authentication, set the SEC\_CASE\_SENSITIVE\_LOGON initialization parameter. Only users who have the SYSDBA or SYSOPER privilege can set the SEC\_CASE\_SENSITIVE\_LOGIN parameter. Set it to TRUE to enable case sensitivity or FALSE to disable case sensitivity.

For greater security, Oracle recommends that you enable case sensitivity in passwords. However, if you have compatibility issues with your applications, you can use this parameter to disable password case sensitivity. Examples of application compatibility issues are passwords for your applications being hard-coded to be case insensitive, or different application modules being inconsistent about case sensitivity when sending credentials to start a database session.

## Password File

---

Passwords stored in the password file have also become case sensitive and store the same kind of hash value in addition to the old-style verifier, this is reflected by the ignorecase argument to the orapwd utility, it is 'n' (no) by default so when you create a password file for remote sysdba and sysoper authentication, the SYS password will be case sensitive.

Naturally, when you grant SYSDBA to an existing user, check if this user has the 11G type password listed in DBA\_USERS.PASSWORD\_VERSIONS otherwise Oracle can only store the old-style verifier for that user in the passwordfile and the password will then not be case sensitive.

If the password of the user that was granted SYSDBA or SYSOPER is changed later, Oracle will still store the new verifier and the password will then become case sensitive.