



Glossary

accent-insensitive sort An accent-insensitive sort is a new sorting option in Oracle Database 10g. You can use the session parameter `NLS_SORT` to specify the linguistic sort name. The `_AI` value is suffixed to a sort name to indicate an accent-insensitive sort.

Active Session History (ASH) Active Session History (ASH) consists of recent session activity. ASH is actually a rolling buffer in memory, and you can view the statistics using the `V$ACTIVE_SESSION_HISTORY` view. The Manageability Monitor (MMON) process flushes ASH data to disk every 60 minutes.

ADDM See Automatic Database Diagnostic Monitor.

advisory framework The advisory framework consists of a set of server-based advisors that are designed around a common infrastructure. Automatic Database Diagnostic Monitor (ADDM) is part of the advisory framework, as are the SQL Tuning Advisor and other advisors. All the advisors have a uniform interface and a common data source: the Automatic Workload Repository (AWR).

ALERT_QUEUE You can subscribe to the `ALERT_QUEUE` and read the `ALERT_QUEUE` if you wish to create your own tool to display database alerts, rather than using the Oracle Database 10g Database Control.

ALL_ROWS `ALL_ROWS` is the default value for the `OPTIMIZER_MODE` initialization parameter. The optimizer uses a cost-based approach for all SQL statements in the session, regardless of the presence of statistics, and optimizes with a goal of best throughput (minimum resource use to complete the entire statement).

ARBn `ARBn` is an Automatic Storage Management (ASM) background process that performs the rebalancing of data extents. You can have multiple ARB processes, named `ARB1`, `ARB2`, and so on.

ASH See Active Session History.

ASM See Automatic Storage Management.

ASM_DISK_GROUPS The `ASM_DISK_GROUPS` initialization parameter specifies the list of names of disk groups to be mounted by an ASM instance at instance startup or when you issue the `ALTER DISKGROUP ALL MOUNT` statement.

ASM_DISKSTRING The `ASM_DISKSTRING` initialization parameter specifies an operating system-dependent value used by ASM to limit the number of disks considered for discovery. When a new disk is added to a disk group, the ASM instance will discover the new disk using the value of `ASM_DISKSTRING`. If you omit this parameter, the ASM instance will find all disks to which it has read/write access.

ASM file aliases Automatic Storage Management (ASM) alias filenames provide a more intuitive way of referring to ASM files, rather than using the fully qualified names (system aliases) that ASM generates when it creates new files.

ASM file templates Automatic Storage Management (ASM) files are named collections of attributes that you apply to a file during its creation. Templates simplify file creation. Templates are applied to a file, but are actually associated with a disk group.

ASM filenames Automatic Storage Management (ASM) filenames can be of the following forms: fully qualified, numeric, alias, alias with template, incomplete, and incomplete with template. The correct filenameing form depends on the context in which you intend to use the file. The three different contexts are referencing an existing file, creating a single new file, and creating multiple new files.

ASM files Automatic Storage Management (ASM) files are Oracle database files stored in ASM disk groups. All ASM files are Oracle Managed Files (OMF), and you must back up them with the Recovery Manager (RMAN). ASM doesn't manage binary files, alert logs, trace files, or password files.

ASM instance An Automatic Storage Management (ASM) instance is the special instance you need to start in order to use ASM in your database. ASM instances are unique in the sense that they don't mount databases, but merely make ASM files available to regular database instances. Databases communicate with an ASM instance to get information about the ASM file layouts.

ASM_POWER_LIMIT The `ASM_POWER_LIMIT` initialization parameter pertains to an ASM instance, and it controls the speed of a disk-rebalancing operation. The range is 1 to 11, with 1 as the default and 11 providing the fastest rebalancing when you add or remove disks from an ASM disk storage system. The higher the limit, the faster rebalancing will complete. Lower values will take longer, but they consume fewer processing and I/O resources.

4 Glossary

ASMB ASMB is a background process that exists in each database that uses Automatic Storage Management (ASM). When you start the database, the ASMB process connects to the ASM process and a foreground process, and conducts all communication between the database and the ASM instance.

ATO See Automatic Tuning Optimizer.

AUDIT_TRAIL The `AUDIT_TRAIL` initialization parameter enables or disables database auditing. `AUDIT_TRAIL` can take the values `db` or `true`, `os`, and `db_extended`. The `db_extended` value enables database auditing and sends audit records to the database audit trail (the `SYS.AUD$` table). It also populates the `SQLBIND` and `SQLTEXT CLOB` columns of the `SYS.AUD$` table.

automatic channel failover When you backup your database with the Recovery Manager (RMAN), one of the channels may fail during the process. Automatic channel failover is the new feature whereby RMAN will attempt to complete the backup on another channel.

automatic checkpoint tuning Oracle Database 10g automatically tunes checkpointing activity to minimize recovery time while reducing the impact of checkpointing on performance. You don't need to set any checkpoint-related parameters. You can either enable automatic checkpoint tuning by setting a nonzero value for the `FAST_START_MTTR_TARGET` initialization parameter or omit it from your `init.ora` or `SPFILE` file.

Automatic Database Diagnostic Monitor (ADDM) The Automatic Database Diagnostic Monitor (ADDM) is a self-diagnostic engine that is part of Oracle Database 10g. ADDM helps the database diagnose its performance and determine how to resolve problems. ADDM runs automatically after each Automatic Workload Repository (AWR) statistics capture and performs analysis on a proactive basis. ADDM diagnoses the performance of the system, using time as a common currency between components.

automatic maintenance tasks The automatic maintenance tasks infrastructure enables Oracle Database 10g to automatically perform routine administrative tasks such as collecting optimizer statistics during a predefined maintenance window. By default, the maintenance window starts at 10 P.M. every night and lasts until 6 A.M. the next morning, and also includes the entire weekend. You can customize attributes of the maintenance window, including start and end time, frequency, and days of the week.

automatic optimizer statistics collection Oracle Database 10g fully automates optimizer statistics collection. The Oracle Scheduler comes with a built-in job called `GATHER_STATS_JOB`, which executes the `DBMS_STATS.GATHER_DATABASE_STATS_JOB_PROC` procedure to automatically collect the statistics. All you need to do is ensure that the `STATISTICS_LEVEL` initialization parameter is set to `TYPICAL` or `ALL` (not `BASIC`). Oracle automatically gathers statistics for all database objects that have stale or missing statistics.

Automatic Shared Memory Management Automatic Shared Memory Management simplifies System Global Area (SGA) memory management by automatically allocating memory among the four auto-tuned parameters: database buffer cache (default pool), shared pool, large pool, and Java pool. To use Automatic Shared Memory Management, set the `SGA_TARGET` initialization parameter to a nonzero value and set the `STATISTICS_LEVEL` initialization parameter to `TYPICAL` or `ALL`.

automatic shared server configuration The automatic shared server configuration helps you to easily switch between shared server and dedicated server configurations without needing to start a dispatcher.

Automatic Storage Management (ASM) Automatic Storage Management (ASM) is an integration of the file system and a built-in volume manager for Oracle database files. It extends the concept of stripe and mirror everything (SAME) principles to optimize disk performance. ASM simplifies the optimal layout of datafiles, control files, and log files. Database files are automatically distributed across all available disks, and database storage is rebalanced whenever storage configuration changes.

Automatic Tuning Optimizer (ATO) The SQL Tuning Advisor calls the Automatic Tuning Optimizer (ATO) to perform four types of analysis: statistics analysis, SQL profiling, access path analysis, and SQL structure analysis.

Automatic Undo Management Automatic Undo Management is the recommended way to manage the undo in Oracle databases, although you could still do it manually. Once you specify the automatic undo management mode in the `init.ora` (or `spfile`) file and create an undo tablespace, Oracle will automatically create and manage the necessary undo segments.

automatic undo retention tuning Oracle Database 10g automatically tunes the initialization parameter `UNDO_RETENTION` to control retention of undo information in rollback segments.

Automatic Workload Respository (AWR) The Automatic Workload Repository (AWR) is an embedded, self-managing repository that captures workload information and performance-related statistics on a routine basis. The database uses information contained in the repository for self-management activities.

AWR See Automatic Workload Repository.

backup compression If disk space is an issue, or your media-management software does not support compression, you can use the Recovery Manager (RMAN) backup compression feature to compress RMAN backup sets.

base statistics Base statistics refer to raw performance statistics, such as the number of physical reads since instance startup. Metrics are derived from base statistics.

bigfile tablespace A bigfile tablespace is a large tablespace containing just one file. Oracle Database 10g lets you create both bigfile tablespaces and traditional tablespaces, now called smallfile tablespaces. The default tablespace type is smallfile. The bigfile tablespace type makes a tablespace logically equal to a datafile.

block change tracking Oracle tracks the physical location of database changes in a new file called the change-tracking file. The Recovery Manager (RMAN) uses the change-tracking file to determine which data blocks to read during an incremental backup, thus making incremental backups faster by avoiding the reading of entire datafiles.

calendaring expression A calendaring expression is the primary method of setting the `REPEAT_INTERVAL` attribute for a Scheduler job. A calendaring expression has three attributes: frequency, repeat interval, and an optional specifier.

case-insensitive sort A case-insensitive sort is a new sorting option in Oracle Database 10g. You can use the session parameter `NLS_SORT` to specify the linguistic sort name. The `_CI` value is suffixed to a sort name to indicate a case-insensitive sort.

CATALOG START WITH The `CATALOG START WITH` command in the Recovery Manager (RMAN) will look at all files that are not a part of the catalog already, starting with the specified directory. This saves you time and effort in remembering cryptic filenames.

change-aware incremental backups By using the new change-tracking file, the Recovery Manager (RMAN) can avoid scanning the entire datafile during an incremental backup. Instead, it will scan only the changed blocks of data, thus reducing backup time.

change-tracking file Oracle tracks the physical location of database changes in a new file called the change-tracking file. The Recovery Manager (RMAN) uses the change-tracking file to determine which data blocks to read during an incremental backup, thus making incremental backups faster by avoiding the reading of entire datafiles.

Clone Database Tool The Clone Database Tool is an Enterprise Manager wizard that lets you step through a database cloning operation. With this tool, which uses RMAN to perform the cloning, you can clone a database to another Oracle Home while the source database is open.

column-level VPD Virtual Private Database (VPD) is now more fine-grained. Column-level VPD enables you to enforce row-level security when a security-relevant column is referenced in a query. You can now enforce VPD rewrite when a query references a particular column. You thus enforce access control based on the security columns that are accessed by a user.

column-masking behavior By default, when a query references a security-relevant column, column-level VPD restricts the number of rows returned. By using column-masking behavior (using the `SEC_RELEVANT_COLS_OPT` parameter of the `DBMS_RLS.ADD_POLICY` procedure), all rows will be displayed, including those that reference security-relevant columns.

COMPATIBLE The `COMPATIBLE` initialization parameter allows you to use a new release of Oracle, while at the same time guaranteeing backward-compatibility with an earlier version. The parameter specifies the release with which Oracle must maintain compatibility.

compressed backups You can now use binary compression to reduce the size of your backup sets. You must have the compatibility set to at least 10.0.0 and use the following Recovery Manager (RMAN) command:

```
BACKUP AS COMPRESSED BACKUPSET DATABASE;
```

configurationless client connect Configurationless client connect is a new connection identifier for SQL clients, which doesn't depend on any configuration files and doesn't require name lookup services. In the client identifier syntax, all you must have are the host, port, and service name to connect to a database server. This is also known as the easy naming connect method. This method simplifies network management by enabling a database connection to be made without the need to first configure a network service name or perform other network configuration on the client. To configure the easy connect method, you specify the keyword `EZCONNECT` in your `sqlnet.ora` file.

context-sensitive policies A context-sensitive policy changes when certain session context attributes change within a user's session. The database reevaluates the policy function each time there is a context change. The policy predicates will thus change, based on session context changes.

cross-platform transportable tablespaces Cross-platform transportable tablespaces are tablespaces that can be transported across different operating system platforms.

CTWR The CTWR is the new Oracle background process that writes database change information to the change-tracking file.

Data Pump The Data Pump export and import utilities provide high-speed movement of data from one database to another. These utilities offer several advantages over the original Oracle export and import utilities, including the ability to completely restart export and import jobs, detach and reattach to long-running jobs, estimate how much disk space an export job would consume, and perform export and import operations over the network. The Data Pump utilities also support fine-grained object selection.

DATA_PUMP_DIR `DATA_PUMP_DIR` is an environment variable that you can define to specify the directory object name in Data Pump. When you use this variable, you don't need to specify the `DIRECTORY` parameter in a Data Pump job.

Data Pump direct path Direct path is one of two access methods that Data Pump supports to load and unload table data. If it is possible, Data Pump tries to use the direct-path access method to load and unload data. The other access method is external tables.

Database Control (Enterprise Manager) Database Control is the new name for the Oracle Enterprise Manager GUI for managing Oracle databases. You can use the Database Control for managing single databases, and the software comes with the Oracle Database 10g Server software.

database features usage tracking Oracle Database 10g automatically tracks usage of various database features. This enables you to evaluate how each feature is being utilized in the database.

Database Upgrade Assistant (DBUA) The Database Upgrade Assistant (DBUA) is a GUI-based Oracle tool that simplifies the task of upgrading to Oracle Database 10g from an older version of Oracle software.

database wait model The database wait model enables you to determine what sessions are waiting for, keeps a history of waits and wait durations per session, and maintains wait statistics for SQL statements in a dynamic performance view.

datafile conversion If you are transporting tablespaces across platforms, the files in both platforms must have the same endian format. If the endian formats are different, you need to convert the datafiles either before or after you move the datafiles that belong to the tablespaces you're transporting.

DB_FLASHBACK_RETENTION_TARGET The `DB_FLASHBACK_RETENTION_TARGET` initialization parameter specifies the upper limit (in minutes) on how far back in time the database may be flashed back. Of course, how far back you can flashback a database depends on how much flashback data Oracle has kept in the flash recovery area.

DB_RECOVERY_FILE_DEST The `DB_RECOVERY_FILE_DEST` initialization parameter specifies the default location for the flash recovery area.

DB_RECOVERY_FILE_DEST_SIZE The `DB_RECOVERY_FILE_DEST_SIZE` initialization parameter specifies (in bytes) the limit on the total space allocated for the database recovery files created in the flash recovery area.

DBA_DATAPUMP_JOBS The `DBA_DATAPUMP_JOBS` view displays all Data Pump jobs in the database.

DBA_ENABLED_TRACES The `DBA_ENABLED_TRACES` view displays information about enabled SQL traces. The view shows the following types of traces: `CLIENT_ID`, `SERVICE`, `SERVICE_MODULE`, and `SERVICE_MODULE_ACTION`.

DBA_OUTSTANDING_ALERTS The `DBA_OUTSTANDING_ALERTS` view describes alerts that have yet to be resolved.

DBA_SCHEDULER_JOB_LOG The `DBA_SCHEDULER_JOB_LOG` view displays log information for all Scheduler jobs in the database.

DBA_SCHEDULER_JOB_RUN_DETAILS The `DBA_SCHEDULER_JOB_RUN_DETAILS` view displays log run details for all Scheduler jobs in the database.

DBA_SERVER_REGISTRY The `DBA_SERVER_REGISTRY` view displays information about the components loaded into the database. It helps you determine the success of your database upgrade process.

DBMS_ADVISOR `DBMS_ADVISOR` is the database package that is the foundation for the Server Manageability Suite of advisors (advisory framework) that identify and resolve performance problems related to the various database server components.

DBMS_AQADM The `DBMS_AQADM` package provides procedures to manage Oracle Streams Advanced Queuing (AQ) configuration and administration information.

DBMS_DATAPUMP The `DBMS_DATAPUMP` package is used to perform the Data Pump export and import of data.

DBMS_METADATA The `DBMS_METADATA` package lets you retrieve metadata from the database dictionary.

DBMS_MONITOR The `DBMS_MONITOR` package lets you use PL/SQL for controlling additional tracing and statistics gathering.

DBMS_MVIEW The `DBMS_MVIEW` package enables you to understand the capabilities of materialized views and potential materialized views, including their rewrite capability. The package also enables you to refresh materialized views and purge logs.

DBMS_RLS The `DBMS_RLS` package contains the procedures to implement and manage fine-grained access control, which is used to implement Virtual Private Database (VPD).

DBMS_SCHEDULER The `DBMS_SCHEDULER` package provides a collection of functions and procedures that help you manage the Oracle Scheduler facility.

DBMS_SERVER_ALERT The `DBMS_SERVER_ALERT` package lets you issue alerts when some threshold has been violated.

DBMS_SQLTUNE The `DBMS_SQLTUNE` package provides the interface to tuning SQL statements.

DBUA See Database Upgrade Assistant.

dense data Dense data is data where there is a row for each combination of the various dimensions.

disk group An Automatic Storage Management (ASM) disk group is a set of disks managed as a single logical unit. You add or remove storage from a disk group by adding or removing ASM disks from a disk group. All ASM files are stored in a disk group.

disk group mirroring Automatic Storage Management (ASM) disk groups are mirrored at the extent level to protect data. There are three types of mirroring: external redundancy, normal redundancy, and high redundancy.

disk group rebalancing Automatic Storage Management (ASM) automatically dynamically rebalances disk groups when you add or remove disks from a disk group. ASM always keeps I/O balanced by moving around only the amount of data proportional to the amount of storage that is being added or removed from the disk group.

DROP DATABASE The new `DROP DATABASE` command drops a database by deleting all database files, all online logs, control files, and server parameter files (SPFILES).

DURATION `DURATION` is a Recovery Manager (RMAN) backup command option that can take the values `PARTIAL`, `MINIMIZE TIME`, and `MINIMIZE LOAD`. The `DURATION` option lets you control the speed and duration of your backup operations.

dynamic policies Dynamic policies apply to the Virtual Private Database (VPD) concept. Dynamic policies change based on certain variables such as the time of day, thus returning different sets of rows for the same query. Previously, all VPD policies were dynamic by default, which meant that Oracle executed the function for each DML statement. You can now use new static policies.

end-to-end application tracing The new end-to-end application tracing capability simplifies the debugging of performance problems in multitier environments. It also makes debugging and tracing of applications, which use `MODULE` and `ACTION` name annotation, easier and more efficient.

endian format Endian format is the byte ordering of files. When you transport tablespaces across platforms, you must ensure that the endian format is identical on both platforms. If not, you need to convert the datafiles to ensure that they are the same on the source and target platforms.

ESTIMATE The `ESTIMATE` parameter is a Data Pump export parameter and specifies the method that export will use to estimate how much disk space each table in the export job will consume (in bytes).

ESTIMATE_ONLY The `ESTIMATE_ONLY` Data Pump parameter estimates the space that an export or import job would require, without actually performing the export operation.

EXCLUDE The `EXCLUDE` Data Pump parameter enables you to exclude a database object from an export or import operation.

EXPLAIN_MVIEW The `EXPLAIN_MVIEW` procedure of the `DBMS_MVIEW` package tells you whether a materialized view is fast refreshable or eligible for general query rewrite.

external table access driver Data Pump provides an external table access driver (`ORACLE_DATAPUMP`) that reads and writes files to enable high-speed loading and unloading of database tables.

external table population External table population is the ability to load data into platform-independent, Oracle proprietary flat files from a database table.

external tables External tables enable you to use external data as a virtual table that you can view and join to other tables, without needing to load the data in actual database tables. External tables are read-only (no DML activity is possible).

EZCONNECT See configurationless client connect.

failure group A failure group is an Automatic Storage Management (ASM) concept that refers to the set of disks inside a particular disk group that share a common resource, whose failure must be tolerated. A common example is a set of disks that connect to a common SCSI controller.

fast incremental backups The fast incremental backups feature speeds up incremental backups by tracking block changes in the new change-tracking file. The new background process CTWR writes changed block information to the change-tracking file, and the Recovery Manager (RMAN) will back up only the changed blocks.

fast ramp-up Fast ramp-up is the new feature wherein the database uses Automatic Workload Repository (AWR) statistics to decide the number of undo segments to bring online when you start an instance or when you switch an undo tablespace.

fine-grained access control Fine-grained access control enables you to build applications that enforce security policies at a low level of granularity. (These policies are also referred to as VPD policies.) You can use it, for example, to restrict an employee to see only his or her own information in a database table.

FIRST_ROWS You can specify `FIRST_ROWS` as the value for the `OPTIMIZER_MODE` initialization parameter. Under the `FIRST_ROWS` setting, Oracle will use a mix of cost and heuristics to find a best plan for fast delivery of the first few rows.

flash recovery area The flash recovery area is a unified storage location for backups, archive logs, and any other files needed for Oracle recovery. The backup files are automatically deleted after they are backed up by the Recovery Manager (RMAN). The flash recovery area is also the location for saving the Flashback Database logs.

Flashback Database The Flashback Database feature lets you quickly bring your database to a prior point in time by undoing all the changes that have taken place since that time. You enable the Flashback Database feature after you first configure a flash recovery area, and then set the retention target with the `DB_FLASHBACK_RETENTION_TARGET` initialization parameter. The flashback time interval will actually depend on the number of flashback database logs stored in the flash recovery area.

Flashback Database logs Flashback Database logs are a new type of database log files, which are analogous to redo logs. When you turn on the Flashback Database feature, Oracle will start logging before images of data blocks that are being changed. These before images are saved in the Flashback Database logs, which are automatically created and managed in the flash recovery area. The new background process RVWR (recovery writer) writes the Flashback Database data from the flashback buffer to the Flashback Database logs.

Flashback Drop The Flashback Drop feature lets you undo the effects of a `DROP TABLE` statement. The dropped tables are saved in the Recycle Bin. The command you actually use to flashback a table is of the format `FLASHBACK TABLE table-name TO BEFORE DROP`.

Flashback Query The Flashback Query feature lets you query a database as of a certain clock time or a system change number.

Flashback Table The Flashback Table feature lets you recover a table to a point in time in the past without restoring a backup.

Flashback Transaction Query The Flashback Transaction Query capability enables you to examine changes to the database at the transaction level. Using the transaction history, you can diagnose problems, perform analysis, and audit transactions.

FLASHBACK_TRANSACTION_QUERY The `FLASHBACK_TRANSACTION_QUERY` view contains all the SQL statements that help you undo the changes made by specific transactions, using the Flashback Transaction Query feature.

Flashback Versions Query The Flashback Versions Query feature lets you view changes to one or more rows, along with all the metadata of the changes.

flushing the buffer cache Oracle Database 10g lets you flush the buffer cache by using the following command:

```
ALTER SYSTEM FLUSH BUFFER_CACHE;
```

global scripts Global scripts, an enhanced Recovery Manager (RMAN) feature, are scripts that can be used by any database connecting to the target database.

grid computing Grid computing is the coordinated use of a large number of servers and storage acting as one computer. Grid computing lets you automatically provision for spikes on demand, and you'll have computing power available when you need it. Grids are built with low-cost modular components, so you can start small and preserve your investment as your business needs grow.

Grid Control The Grid Control is the centralized version of Oracle Enterprise Manager, which lets you manage several databases and hosts simultaneously. You must install the Grid Control separately from the database.

guaranteed undo retention When you enable guaranteed undo retention by using the `RETENTION GUARANTEE` clause either during the creation of the undo tablespace or by altering it later, the database never overwrites unexpired undo data.

hash-partitioned global indexes In a hash-partitioned global index, each partition contains the values determined by a hash function, which is based on the partitioning key and the number of partitions. These indexes improve performance by spreading contention when the index is monotonically growing, because most of the index insertions occur only on the right edge of an index.

idle timeouts You can set up maximum idle timeouts for a resource plan using the `DBMS_RESOURCE_MANAGER` package. When a session exceeds the `MAX_IDLE_TIME` parameter value specified in the `CREATE_PLAN_DIRECTIVE` procedure, the Process Monitor (PMON) will kill the session.

INCLUDE The `INCLUDE Data Pump` parameter lets you require that only specified object types and objects be included in an export or import job.

INSTANCE_TYPE The `INSTANCE_TYPE` initialization parameter is used in connection with an Automatic Storage Management (ASM) instance to specify whether the instance is a database instance or an ASM instance.

job A job is a user-defined task that is scheduled to run one or more times. A job contains information about the task (the action) to be run and when it should run (the schedule). You can specify the action and schedule when creating the job, or the job can refer to an existing program and schedule.

job class A job class is a group of jobs that have similar characteristics and behavior. A job class is a way of grouping jobs into larger entities, thus enabling you to prioritize among the job classes.

job priority Job priority helps you prioritize the jobs in your database. There are two levels at which you can prioritize a job: at the class level, using resource plans, and within the class, using the job priority attribute of the job. The default priority of a job is 3.

locally managed tablespaces Locally managed tablespaces track all extent information in the tablespace itself by using bitmaps. The `DBMS_SPACE_ADMIN` package provides procedures for administering locally managed tablespaces.

locking statistics You can lock representative statistics when a table's data keeps changing continuously. You can lock a table with or without any statistics.

logging level You can set the logging level of the Scheduler by using the `LOGGING_LEVEL` parameter. You can set this parameter to either `DBMS_LOGGING_FULL` or `DBMS_SCHEDULER.LOGGING_RUNS` at either the job class or the individual job level.

MANAGE SCHEDULER `MANAGE SCHEDULER` is a Scheduler-associated system privilege, which lets a user perform all Scheduler administration tasks. The `SCHEDULER_ADMIN` role has all the privileges of the `MANAGE SCHEDULER` privilege along with the `WITH ADMIN OPTION` clause. The `DBA` role includes the `SCHEDULER_ADMIN` role by default.

Manageability Monitor (MMON) The Manageability Monitor (MMON) background process performs various manageability-related background tasks, including issuing alerts whenever a given metric violates its threshold value and taking snapshots for the Automatic Workload Repository (AWR) feature. The MMON wakes up every minute to compute metric values. The MMON also verifies the thresholds for all the metrics you define and generates the alerts, if necessary.

materialized view decomposition Materialized view decomposition refers to the breaking up of a nonrefreshable materialized view into submaterialized views referenced by the original materialized view. The automatic materialized view decomposition mechanism handles cases where the materialized view-defining query has subqueries, set operators, or inline views. Each of the submaterialized views will be refreshable.

Memory Manager (MMAN) The new Memory Manager (MMAN) background process manages automatic shared memory. MMAN serves as the System Global Area (SGA) memory broker and coordinates the sizing of the individual memory components. It tracks the size of the memory components and monitors memory-resizing operations.

MERGE The MERGE command performs conditional update and insert operations. You may also use an optional DELETE clause with a MERGE command.

metrics Metrics are secondary statistics that the database derives from base statistics. Metrics track the rate of change of a variable. An example of a database metric is the number of physical reads in the database in the past 60 minutes.

MMAN See Memory Manager.

MMNL The MMNL (Manageability Monitor Process Light; shown as the Manageability Monitor Process 2 in a query of the V\$BGPROCESS view) performs frequent and lightweight manageability-related tasks, such as session history capture and metrics computation. If the Automatic Session History (ASH) buffer fills before MMON flushes it, MMNL will flush the data to disk.

MMON See Manageability Monitor.

MODEL See SQL MODEL clause.

network mode export By using the NETWORK_LINK export parameter, you can perform a network export, provided you also specify the name of a valid database link. A network export moves data from a remote database to a dump file set local to the instance running the Data Pump job.

network mode import By using the `NETWORK_LINK` import parameter, you can initiate a network import. The `impdp` client will initiate the import request, and the server contacts the remote source database referenced by `source_database_link`. The server retrieves and writes the data directly to the target database tables. Since the entire import is directly via the network, dump files aren't necessary.

nonthreshold alerts Some server-generated alerts refer to specific database events such as ORA-1555 (snapshot too old) errors. These are not based on any thresholds, so they are called nonthreshold alerts, or stateless alerts. These alerts go directly to the history table, without ever appearing in the `DBA_OUTSTANDING_ALERTS` view.

online segment shrink The online segment shrink feature shrinks segments online and in-place (tables, indexes, and materialized views) that have free space in them, thereby enhancing efficiency of space utilization.

Partition Change Tracking (PCT) Partition Change Tracking (PCT) is the ability to identify which rows in a materialized view are affected by a detail table partition.

partition outer join Partitioned outer joins are extensions to the outer join syntax, and they are used primarily to improve performance and simplify SQL queries for time-based calculations.

PCT See Partition Change Tracking.

PCT refresh When a materialized view is partitioned on the partitioning key column or join-dependent expressions of the detail table, it is better to truncate one or more partitions during a materialized view refresh, and then populate the partition with new data. This variant of the fast refresh is known as a PCT refresh.

PLSQL_OPTIMIZE_LEVEL `PLSQL_OPTIMIZE_LEVEL` specifies the optimization level that Oracle will use to compile PL/SQL library units. The higher the setting of this parameter, the more effort the compiler makes to optimize PL/SQL library units. Oracle claims that setting this parameter to 1 or 2 will make PL/SQL programs run considerably faster.

policy framework The policy-based database framework lets you monitor all targets to see if their configuration conforms to recommended policy configuration settings. These policies are based on Oracle's best practice configuration recommendations. The database configuration framework is built on top of the configuration and metric collection facility of the Oracle Enterprise Manager.

policy function Fine-grained access control enables you to implement security policies with policy functions and to associate these policies with tables or views. The security policy function generates a *WHERE* condition that is appended to relevant SQL statements, thereby restricting user access to rows of data in the table or view.

Post-Upgrade Status Tool The Post-Upgrade Status Tool is invoked by running the `utlu101s.sql` script located in the `$ORACLE_HOME/rdbms/admin` directory, after you upgrade a database to the Oracle Database 10g release. The script will tell you if the individual components of the new database have been successfully upgraded.

POWER The *POWER* clause is used when you rebalance an Automatic Storage Management (ASM) disk group and specifies the power level of a disk rebalancing job. By using the *POWER* clause in a rebalance disk group command, you can override the setting of the default speed set by using the initialization parameter `ASM_POWER_LIMIT`.

proactive tablespace management The Tablespace Advisor gathers and reports historical data about each tablespace's disk space usage and notifies administrators when tablespaces run low on available space.

program A program is a collection of metadata about what task the Scheduler will run.

projected columns When dealing with external table files, which contain rows of data that may be rejected due to errors, the projected column feature enables you to get a consistent result set, independent of the columns referenced by the SQL statement accessing the data. Prior to Oracle Database 10g, only the columns referenced by the SQL statement were projected out by the access driver, and some rows were rejected due to conversion errors or data format errors.

PURGE You need to use the `PURGE` clause when you drop a database table if you want Oracle to remove the table permanently. Otherwise, Oracle will save the table in the Recycle Bin. Using the `DROP TABLE PURGE` command is equivalent to using the old `DROP TABLE` command.

RATIO `RATIO` refers to the new resource allocation method for specifying how much CPU each consumer group or subplan gets, when you use the Database Resource Manager. The `RATIO` method for specifying CPU allocation is for single-level plans that use ratios to specify how CPU is distributed among consumer groups. `EMPHASIS`, the default method, is for multilevel plans that use percentages to specify how CPU is distributed.

RBAL The `RBAL` background process coordinates rebalancing activity for disk groups in an Automatic Storage Management (ASM) instance.

rebalance Rebalance is an Automatic Storage Management (ASM) term that refers to the automatic rebalancing of disk groups when you add or remove disks. ASM performs the rebalancing to maintain balanced I/O across all the disks in a disk group.

Recycle Bin The Recycle Bin is actually a data dictionary table containing information about dropped objects. Dropped tables and any associated objects—such as indexes, constraints, and nested tables—are renamed. You can use the `FLASHBACK TABLE ... TO BEFORE DROP` command to recover a dropped table from the Recycle Bin.

Redo Logfile Size Advisor The Redo Logfile Size Advisor recommends the optimal size of redo log files.

remapping Remapping is a Data Pump concept that refers to the ability to apply transformation during an import operation. You can remap datafiles using the `REMAP_DATAFILE` parameter, tablespaces using the `REMAP_TABLESPACE` parameter, and schemas using the `REMAP_SCHEMAS` parameter.

renaming tablespaces Oracle Database 10g provides the ability to rename tablespaces. You no longer need to create a new tablespace, copy the contents from the old tablespace, and drop the old tablespace.

RESUMABLE_TIMEOUT The `RESUMABLE_TIMEOUT` initialization parameter enables or disables resumable statements. You can also use this parameter to specify a resumable timeout at the system level.

REWRITE_OR_ERROR The `REWRITE_OR_ERROR` hint will issue an error if a query fails to rewrite. Using the `REWRITE_OR_ERROR` hint in a query causes the following error if the query failed to rewrite:

```
ORA-30393: a query block in the statement did not rewrite
```

RUN_TO_COMPLETION `RUN_TO_COMPLETION` is a new value that you can use for the `CPU_MTH` parameter when you create a resource consumer group. It specifies that sessions with the largest active time are scheduled ahead of other sessions. The `CPU_MTH` parameter specifies the distribution of CPU among sessions in a consumer group. The default value is `ROUND_ROBIN`, which uses a round-robin schedule to ensure that sessions are executed fairly.

schedule A schedule indicates when a job should run. It has a start date that specifies the date and time when the schedule starts, an end date that indicates the date and time when the schedule expires, as well as a repeat interval that indicates how often it will repeat.

Scheduler The Scheduler is Oracle's new scheduling feature that gives you the ability to schedule a job to run at a particular date and time. It also provides you the ability to create libraries of sharable Scheduler objects. You can group jobs that share common characteristics into larger entities called job classes, which can be prioritized by controlling the system resources allocated to each of them.

Segment Advisor The Segment Advisor gives advice on whether an object is a good candidate for the new online shrink operation, based on the amount of fragmentation within segments. The Segment Advisor also reports on the historical growth trend of segments.

segment resource estimation The new segment resource estimation feature facilitates space management by giving a size estimate of the size of a new table or index, so you can ensure that sufficient space is available.

segment shrinking Segment shrinking is the ability to shrink a segment in order to eliminate fragmentation within the segment. Shrinking a segment can also lower the high-water mark (HWM) of space usage for that object.

server-generated alerts Server-generated alerts are alerts that the database will proactively send to administrators when a problem is anticipated or any of the user-selected metrics exceeds a defined threshold.

Server Manageability Suite The Server Manageability suite of Advisors is a set of expert systems that helps you identify and resolve various performance problems in your database. These expert systems are the SQL Tuning Advisor, the SQL Access Advisor, the Undo Advisor, the Segment Advisor, and the various memory advisors.

SGA_TARGET The `SGA_TARGET` initialization parameter specifies the total size of all SGA components. If `SGA_TARGET` is specified, the following SGA memory pools are automatically sized: buffer cache (`DB_CACHE_SIZE`), shared pool (`SHARED_POOL_SIZE`), large pool (`LARGE_POOL_SIZE`), and Java pool (`JAVA_POOL_SIZE`).

shared policies Shared policies apply to Virtual Private Directory (VPD). They let you apply the same policy to several objects. You can have both shared static and shared context-sensitive policy types.

SKIP_UNUSABLE_INDEXES The `SKIP_UNUSABLE_INDEXES` initialization parameter enables or disables the use and reporting of tables with unusable indexes or index partitions.

smallfile tablespace In Oracle Database 10g, you can have two types of tablespaces: bigfile and smallfile. Smallfile is the default tablespace type, and it is the same as the standard tablespaces used in previous Oracle versions.

snapshot The Automatic Workload Repository (AWR) automatically collects performance data (by default, on an hourly basis). These statistics are known as snapshots. You can also manually create snapshots.

snapshot baseline Baselines are defined on a pair of snapshots, and serve to tag sets of snapshot data to compare performance against. Baselines help you retain snapshot data belonging to a representative period in the past to compare against the current database behavior.

sorted hash cluster Sorted hash clusters are new data structures that let you retrieve data faster in applications where data is requested in the order in which it was inserted.

sparse data Sparse data refers to data with gaps, where no row exists for a certain combination of dimensions.

SQL Access Advisor The SQL Access Advisor identifies and helps resolve performance problems relating to SQL statement execution by recommending which indexes or materialized views to create, drop, or retain.

SQL MODEL clause The SQL MODEL clause provides a powerful building block for complex calculations such as forecasts and budgets. Using the MODEL clause, you can treat relational tables as *n*-dimensional arrays and specify interrow references without complex SQL joins and unions. With the MODEL clause, you can define a multidimensional array on query results, and then apply sophisticated rules on the array to calculate new values. By integrating advanced calculations into the database, performance, scalability, and manageability are enhanced significantly compared to performing the calculations using external products.

SQL profile The Automatic Tuning Optimizer (ATO) creates profiles of SQL statements, which are called SQL profiles, consisting of auxiliary statistics specific to that statement. SQL profiles address the problem of the query optimizer sometimes using incorrect statistics, by collecting additional information using sampling and partial execution techniques to verify and, if necessary, adjust the optimizer estimates.

SQL regular expressions Regular expressions enhance search-and-replace capabilities. This functionality is implemented through new functions that are regular expression extensions to existing Oracle functions such as LIKE, REPLACE, and INSTR.

SQL Tuning Advisor The SQL Tuning Advisor is a new database advisor that eliminates manual tuning of SQL statements. It provides tuning advice in the form of precise actions along with their expected performance benefit. You can use this feature for SQL tuning of packaged applications, without requiring any modification to the application code.

SQL Tuning Set (STS) A SQL Tuning Set (STS) is a new database object used for capturing SQL workload information and includes one or more SQL statements, a list of bind values, and basic execution statistics such as elapsed time, CPU time, and so on. SQL Tuning Sets provide the basic framework for capturing, managing, and tuning a representative SQL workload. SQL Tuning Sets allow selective, on-demand tuning of custom workloads. You can capture any SQL statements of interest and store them in an STS for future tuning without having to build and maintain SQL scripts for tuning purposes.

SQLTUNE_CATEGORY The `SQLTUNE_CATEGORY` initialization parameter specifies the category name that sessions can use to look up SQL profiles during SQL compilation. By default, all profiles are created in the default category. This means that all user sessions where the `SQLTUNE_CATEGORY` initialization parameter is set to `DEFAULT` can use the profile.

STARTUP UPGRADE `STARTUP UPGRADE` is the new way to start a database before you run the upgrade script. This command will start the database in the upgrade mode and sets system initialization parameters to values required to let the database upgrade scripts run. You need to use this mode only when you first start an older version of an Oracle database with a newer version (Oracle Database 10g) of the Oracle Database Server.

static policies Static policies apply to the Virtual Private Database (VPD) concept. Static policies enforce an unchanging policy. Previously, all VPD policies were dynamic by default, which meant that Oracle executed the function for each DML statement. Static policies enforce the same predicate for access control, regardless of the user, thus ensuring the same policy predicate for all users. The new static policies improve performance, since they don't always reexecute policy functions for each DML statement.

STATISTICS_LEVEL The `STATISTICS_LEVEL` initialization parameter specifies the level of collection for database and operating system statistics, which the database needs for many purposes, including making self-management decisions. The default value is `TYPICAL`. The other possible values are `BASIC` and `ALL`. Oracle recommends that you use the default value of `TYPICAL` to get the statistics necessary for most database management tasks, including using the various advisors.

SWITCH DATABASE `SWITCH DATABASE` is a new Recovery Manager (RMAN) backup command that enables you to perform a fast restore of your database. Unlike a normal restore, this command doesn't copy any database files. Instead, it merely points the datafile pointers in the control file to the backup file location and begins recovery.

SYSAUX tablespace `SYSAUX` is a new, mandatory Oracle Database 10g tablespace. It provides a centralized location for all auxiliary database metadata that isn't a part of the `SYSTEM` tablespace and reduces the number of tablespaces created by default.

table monitoring Table monitoring refers to the tracking of changes to tables to determine if statistics collection is necessary. In Oracle Database 10g, table monitoring is enabled by default. The database will collect optimizer statistics automatically by determining which objects have stale statistics or no statistics.

temporary tablespace group A temporary tablespace group lets you use temporary space from multiple tablespaces. Using a tablespace group helps reduce the contention for temporary tablespace when the database is performing huge sorts.

threshold-based alerts Server-generated alerts can be threshold-based or nonthreshold-based. Threshold-based alerts include alerts for events like CPU time per service call, for which you might specify warning and critical thresholds. You can use the `SET_THRESHOLDS` procedure of the `DBMS_SERVER` package to set thresholds for various events.

time model statistics Time model statistics indicate the amount of DB time associated with a process from the `V$SESS_TIME_MODEL` and `V$SYS_TIME_MODEL` views. Time model statistics help the database keep track of the time spent in performing internal operations such as parse, execute, input/output, and so on. The database uses this information for making self-tuning decisions and to diagnose performance problems.

TRANSFORM The `TRANSFORM` parameter applies to Data Pump import. It lets you change the DDL statements for an object, by letting you exclude either or both segment attributes, as well as the storage clauses of objects.

trcsess The `trcsess` command-line utility consolidates information from several trace files into a single file, to enable meaningful, session-based tracing.

TUNE_MVIEW The `TUNE_MVIEW` procedure of the `DBMS_ADVISOR` package shows you how to optimize your materialized view creation statements. It also suggests ways to ensure a fast refresh and general query rewrite of materialized views.

UNCATALOG You use the `UNCATALOG` command in the Recovery Manager (RMAN) to remove backup pieces from the recovery catalog, as shown here:

```
RMAN> change backuppiece 'file_name' UNCATALOG;
```

Undo Advisor The Undo Advisor is part of the advisory framework. It helps database administrators in sizing the undo tablespace and in setting the size of the `UNDO_RETENTION` parameter.

UNDO_RETENTION The `UNDO_RETENTION` parameter specifies (in seconds) the minimum time for which Oracle retains undo information.

uniform audit trail Oracle tracks a set of identical fields for both standard and fine-grained auditing to enable you to easily analyze database activity. This tracking of identical information is called the uniform audit trail. There is a new audit trail view, `DBA_COMMON_AUDIT_TRAIL`, which combines standard and fine-grained audit records.

Upgrade Information Tool The new Upgrade Information Tool, invoked by running the `utlu101i.sql` script, provides information before a database upgrade by performing some preliminary checks on the existing database. For example, it checks whether there is enough space and whether there are any obsolete initialization parameters.

user-specified quoting character Database programmers can now choose any convenient delimiter and define it as the quoting character, using the new quote operators.

UTL_COMPRESS The `UTL_COMPRESS` package provides a set of data compression utilities that have the familiar functionality of the `zip` and `unzip` (`gzip` and `gunzip`) utilities in a PL/SQL environment. This package also handles LOBs and has features to add and later extract several pieces into a single archive.

UTL_MAIL The `UTL_MAIL` package is a utility for managing e-mail. It includes commonly used e-mail features, such as attachments, CC, BCC, and return receipt.

VPD See Virtual Private Database.

Virtual Private Database Virtual Private Database (VPD) enables you to enforce fine-grained security by applying security policies directly to tables, views, or synonyms. When a user accesses a table, view, or synonym that is protected with a VPD policy, the database dynamically modifies the user's SQL statement, using any condition that can be expressed in or returned by a function. You can apply VPD policies to `SELECT`, `INSERT`, `UPDATE`, `INDEX`, and `DELETE` statements.

window A window is a time duration for which you can specify certain resource allocations. Like a schedule, it has a start date when the window becomes active or opens and a duration that specifies how long the window will be open, as well as an end date. It may also have a repeat interval that specifies how often the window will open. You generally associate a resource plan with a window in order to specify how resources should be allocated among the job classes.

window group A window group represents a list of windows and facilitates management of windows. For example, you can combine the weekend and holiday windows into a single downtime window group.

window priority Window priorities are relevant when you are dealing with overlapping windows. If two windows with the same priority overlap, the active window stays open. If two windows with different priorities overlap, the higher priority window will open, and the lower priority window will close.