

New initialization Parameters in Oracle Database 11g

Parameter	Description	Default
Introduced in Database 11g Release 1 (11.1)		
asm_preferred_read_failure_groups	A comma-delimited list of strings that specifies the failure groups that contain preferred read disks. This parameter is used on clustered ASM instances and is instance specific.	NULL
client_result_cache_lag	Specifies the amount of time (in milliseconds) that the client cache can lag behind any changes made to the database that affect the result set. Default is 3 seconds.	3000
client_result_cache_size	Specifies the maximum size (in bytes) of the client result set cache. All OCI clients inherit this settings, but it can be overridden by the OCI_RESULT_CACHE_MAX_SIZE client configuration parameter. The default value (0) set the the result cache to disabled.	0
commit_logging	Advanced parameter that is used to control how redo is batched by Log Writer. If set to IMMEDIATE, the database writes commit records to the disk immediately, generating an I/O operation for every commit. This is the default behavior if the parameter is not set. In BATCH mode the database buffers the redo information, thus several I/O operations are batched. Valid options: IMMEDIATE or BATCH.	No default value
commit_wait	Advanced parameter used to control when the redo for a commit is flushed to the redo logs. This parameter is set on system, session and transaction level. WAIT: The commit records are immediately flushed do disk. Session does not return control until the records are written. NOWAIT: The database doesn't wait for the commit to succeed, but returns control to client right away. FORCE_WAIT: If set at system level, the database ignores session level settings and uses WAIT. If set at session level, the database ignores what is set in a transaction and uses WAIT.	No default value
control_management_pack_access	Specifies which management packs are available and should be active. Valid options: NONE, DIAGNOSTIC, DIAGNOSTIC+TUNING	DIAGNOSTIC+TUNING
db_lost_write_protect	Enables the lost write detection. Lost write occurs when the file system acknowledges that a data block is written, while in fact it still not permanently stored. <u>On primary databases:</u> TYPICAL - the instance logs buffer cache reads in the redo log (for read-write tablespaces); FULL - the instance logs buffer cache reads in the redo log (for read-write and read-only tablespaces); <u>On standby databases (or primary in media recovery):</u> TYPICAL/FULL - the instance performs lost write detection;	NONE

	Valid options: NONE, TYPICAL, FULL	
db_securefile	Specifies if the instance treats LOB files as SecureFiles. NEVER - LOBs that are specified as SecureFiles are created as BasicFile LOBs; PERMITTED - LOBs that are specified as SecureFiles are created as SecureFiles LOBs; ALWAYS - All LOBs are created as SecureFile LOBs; IGNORE - SecureFiles and their options are ignored; Valid options: NEVER, PERMITTED, ALWAYS, IGNORE	PERMITTED
db_ultra_safe	Specifies protection level values for set of other database parameters. <u>OFF</u> : No changes to DB_BLOCK_CHECKING, DB_LOST_WRITE_PROTECT and DB_BLOCK_CHECKSUM, if they are explicitly set; otherwise they are set to default values; <u>DATA_ONLY</u> : DB_BLOCK_CHECKING is set to MEDIUM, DB_LOST_WRITE_PROTECT is set to TYPICAL, DB_BLOCK_CHECKSUM is set to FULL; <u>DATA_AND_INDEX</u> : DB_BLOCK_CHECKING is set to FULL, DB_LOST_WRITE_PROTECT is set to TYPICAL, DB_BLOCK_CHECKSUM is set to FULL;	OFF
ddl_lock_timeout	Specifies a time limit (in seconds) for how long the DDL statements will wait to acquire a lock. The default value is NOWAIT(0). If a lock can not be acquired in the specified amount of time an error is raised (ORA-00054). Valid options: 0 to 1000000	0
diagnostic_dest	Specifies the diagnostics directory (ADR home). The instance generates trace, alert, core and incident files in this directory. Oracle recommends that RAC instances specify the same diagnostic_dest (pointing to a shared disk location).	Derived from \$ORACLE_BASE
global_txn_processes	Specifies the initial number of GTXn processes, that Oracle uses to support global transactions (XA) in RAC environment. There is no need to manually set this parameter, as the database will automatically tune it according to its needs. Valid options: 1 to 20	1
java_jit_enabled	Enables the Java JIT (JustInTime) compiler. Defaults to <i>true</i> for platforms, that support JIT; otherwise its value is <i>false</i> .	OS dependent
ldap_directory_sysauth	Enables (yes) or disables (no) directory-based authentication for users in the SYSDBA and SYSOPER groups. Valid options: yes or no	no
memory_max_target	Specifies the maximum value to which the DBA can increase the memory_target parameter. Valid options: 0 to the size of physical memory available to the database	0
memory_target	Specifies the amount of memory available to the database. The database automatically tunes the used memory to memory_target, allocating and releasing SGA and PGA memory as needed.	0

	Valid options: 152M to memory_max_target	
optimizer_capture_sql_plan_baselines	When set to <i>true</i> , enables the recognition of repeatable SQL statements and generation of SQL plan baselines for such statements.	false
optimizer_use_invisible_indexes	Enables/disables the use of invisible indexes. When set to <i>true</i> , invisible indexes are treated as normal indexes.	false
optimizer_use_pending_statistics	In previous DB versions, the newly gathered optimizer statistics were immediately published. In 11g this behavior can be controlled and statistics can be kept in a pending state until they are published. This parameter specifies whether the optimizer uses statistics in pending state when compiling SQL statements.	false
optimizer_use_sql_plan_baselines	Specifies whether the optimizer will look for a SQL plan baseline for the SQL statement being compiled. If set to <i>true</i> , the optimizer will look for a plan in the SQL Management Base and will cost each of the found plans, choosing the one with the lowest cost.	true
parallel_io_cap_enabled	If set to <i>true</i> the database will cap the default degree of parallelism to no greater than what the I/O system can support. The limits are calculated based on the results of the I/O calibration package (you should run DBMS_RESOURCE_MANAGER.CALIBRATE_IO)	false
plsscope_settings	PL/Scope is a tool that collects data about identifiers in PL/SQL source code at compilation time, and makes this data available in static dictionary views (*_IDENTIFIERS). The identifier data collections is disabled by default. Collection can be set on system, session or per-library unit basis. Valid options: identifiers:none, identifiers:all	identifiers:none
redo_transport_user	The parameter is used to select a different user password for redo transport authentication by setting this parameter to the name of any user who has been granted the SYSOPER privilege. The default is to use the SYS user.	No default value
result_cache_max_result	Specifies the percentage of result_cache_max_size that a single result set can use. Valid options: 0 to 100	5
result_cache_max_size	Sets the maximum amount of SGA memory that can be utilized by the Result Cache. Setting this to 0 disables the Result Cache. Valid options: 0 to OS system dependent value (should be rounded up to the multiple of 32 KB)	Automatically calculated according the SGA size.
result_cache_mode	The default behavior is to add the Result Cache operator only when the query contains hint (/*+ result_cache */). If set to <i>force</i> , the operator is added to the root of all SELECT statements. Valid options: manual, force	manual
result_cache_remote_expiration	Specifies a number of minutes that a cache result based on a remote object will remain valid in the Result Cache. Using the default makes the database to not cache results based on remote objects. Setting this parameter to a non-zero value can produce stale answers (if the remote object has been	0

	changed in the meantime).	
sec_case_sensitive_logon	Specifies whether the login passwords are case sensitive.	true
sec_max_failed_login_attempts	Specifies the number of failed login attempts that the server process tolerates, before it drops the client's connection. Valid options: 1 to unlimited	10
sec_protocol_error_further_action	Specifies how the server process will act upon receiving a bad packet from a client. <u>CONTINUE</u> : The process will continue execution, ignoring the packet <u>(DELAY, number_of_seconds)</u> : The client connections is delayed for a <i>number_of_seconds</i> , before the process accepts its next request. <u>(DROP, bad_packets)</u> : If the number of bad packets reach the <i>bad_packets</i> value, the client connection is terminated. Valid options: CONTINUE, (DELAY, <i>number_of_seconds</i>), (DROP, <i>bad_packets</i>)	continue
sec_protocol_error_trace_action	Specifies the database actions when a bad packet is received from a possibly malicious client. <u>NONE</u> : Bad packets are ignored, no further actions are taken <u>TRACE</u> : A detailed trace file is generated for each bad packet. <u>LOG</u> : A minimal log message is sent to the alert log file. <u>ALERT</u> : An alert is sent via OEM. Valid options: none, trace, log, alert	trace
sec_return_server_release_banner	Specifies if the server returns complete database software information to clients (for OCI connections). Disabling the version information is done by setting this parameter to <i>true</i> .	false
xml_db_events	Enables or disables XML DB Repository events. Valid options: enable, disable	enable
Introduced in Database 11g Release 2 (11.2.0.1)		
deferred_segment_creation	By default segments for non-partitioned tables and their dependent objects (LOBs, indexes) will not be created until the first row is inserted into the table. If set to <i>false</i> , the database will create the appropriate segments together with the database object creation.	true
dst_upgrade_insert_conv	Specifies whether or not internal operators will be allocated on top of TIMESTAMP WITH TIME ZONE (TSTZ) columns of tables which have not been upgraded during the upgrade window of daylight saving time patching for TIMESTAMP WITH TIME ZONE data. When DST_UPGRADE_INSERT_CONV is set to true during the upgrade window of the daylight saving time patching process: SELECT queries on tables with TSTZ data which have not been upgraded will use internal operators on top of TSTZ columns to present TSTZ data as if they were recorded using the new	true

	<p>timezone translation rules.</p> <p>DML on tables with TSTZ data which have not been upgraded will use internal operators on top of TSTZ columns to ensure that the TSTZ data is recorded using the old timezone translation rules in order to be consistent with the existing TSTZ data in the same tables.</p>	
listener_networks	<p>Specifies one or more sets of local & remote listeners for cross-registration. All listeners within the same network_name will cross-register.</p> <p>Valid options:</p> <pre>'((NAME=network_name) (LOCAL_LISTENER=["]listener_address[,...][") [(REMOTE_LISTENER=["]listener_address[,...][")])]' [,...]</pre>	No default value
parallel_degree_limit	<p>Specifies how the optimizer will limit the degree of parallelism, when executing SQL queries, to ensure the parallel processes do not flood the system.</p> <p>CPU: Parallelism is limited by the number of CPUs and is calculated as parallel_threads_per_cpu * cpu_count * number of available instances</p> <p>IO: Parallelism is limited by the I/O capacity. Calculated as the total system throughput divided by the maximum I/O bandwidth per process. DBMS_RESOURCE_MANAGER.CALIBRATE_IO should be run before using this setting.</p> <p>Integer: A numeric value for this parameter specifies the maximum degree of parallelism the optimizer can choose for a SQL statement when automatic degree of parallelism is active. Automatic degree of parallelism is only enabled if parallel_degree_policy is set to <i>auto</i> or <i>limited</i>.</p> <p>Valid options: cpu, io, integer</p>	cpu
parallel_degree_policy	<p>Specifies whether or not automatic degree of Parallelism, statement queuing, and in-memory parallel execution will be enabled.</p> <p>MANUAL: Disables automatic degree of parallelism, statement queuing, and in-memory parallel execution. This reverts the database behavior to 11.2.</p> <p>LIMITED: Statement queuing and in-memory parallel execution is disabled. Automatic degree of parallelism is enabled for statements that access tables or indexes using the PARALLEL clause.</p> <p>AUTO: Automatic degree of parallelism, statement queuing, and in-memory parallel execution is enabled.</p> <p>Valid options: manual, limited, auto</p>	manual
parallel_force_local	<p>Setting the parameter to <i>true</i> restricts the parallel server processes so that they can only operate on the same RAC node where the query coordinator resides (the node on which the SQL statement was executed on).</p>	false
parallel_min_time_threshold	<p>Specifies the minimum execution time a statement should have before the statement is considered</p>	auto

	for automatic degree of parallelism (10 seconds by default). This parameter is only useful for queries that use automatic degree of parallelism. Valid options: auto, integer	
parallel_servers_target	Sets the number of parallel server processes allowed to run parallel statements before statement queuing will be used. When the number of parallel server processes reaches the value of parallel_server_target and parallel_degree_policy is set to <i>auto</i> , the database will start queuing the SQL statements requiring parallel execution. Valid options: 0 to parallel_max_servers	4 * cpu_count * parallel_threads_per_cpu * active_instance_count
Introduced in Database 11g Release 2 (11.2.0.2)		
cursor_bind_capture_destination	Sets the location at which bind variables that are captured from SQL cursors are stored. <u>OFF</u> : No bind variables are captured. <u>MEMORY</u> : Bind variables are captured to memory (V\$ views) <u>MEMORY+DISK</u> : Bind variables are captured to memory and disk (V\$ views, AWR) Valid options: off, memory, memory+disk	memory+disk
db_flash_cache_file	Specifies a filename for the flash memory or a disk group representing a collection of flash memory. This parameter should only be set in combination with db_flash_cache_size. Valid options: <i>filename, diskgroup</i>	No default value
db_flash_cache_size	Sets the Database Smart Flash Cache size. Specified only at instance startup. The Flash Cache is disabled by default (parameter is set to 0). Valid options: 0 to OS dependent	0
db_unrecoverable_scn_tracking	Enables/disables the tracking of unrecoverable (NOLOGGING) direct-path inserts and load operators. Unrecoverable changes can be tracked via the V\$DATAFILE view (UNRECOVERABLE_CHANGE# and UNRECOVERABLE_TIME columns).	true
Introduced in Database 11g Release 2 (11.2.0.3)		
awr_snapshot_time_offset	AWR snapshots normally start at the top of the hour (12:00, 1:00, 2:00, and so on). This parameter allows DBAs to specify an offset (in seconds) for the AWR snapshot start time. 1000000 denotes an automatic mode (the offset is calculated based on the DB name in order to distribute offset times for databases sharing the same server). Valid options: 0 - 3599 or 1000000	No default value

Deprecated Parameters in Database 11g

Parameter	Description
Database 11g Release 1 (11.1)	
background_dump_dest	These parameters are ignored by Oracle Database 11g Release 1, which places trace and core files in a location controlled by the DIAGNOSTIC_DEST initialization parameter.
core_dump_dest	
user_dump_dest	
commit_write	Retained for backward compatibility only. This parameter is replaced by the COMMIT_LOGGING and COMMIT_WAIT parameters.
cursor_space_for_time	Lets you use more space for cursors in order to save time. This parameter is deprecated and is retained for backward compatibility only.
instance_groups	Provides functionality for restricting parallel query operations to a limited number of instances. This parameter is deprecated and is retained for backward compatibility only.
log_archive_local_first	Specifies when the archiver processes (ARCn) transmit redo data to remote standby database destinations. This parameter is deprecated and is retained for backward compatibility only.
plsql_debug	PLSQL_DEBUG specifies whether or not PL/SQL library units will be compiled for debugging. This parameter is deprecated and is retained for backward compatibility only.
plsql_v2_compatibility	PL/SQL Version 2 allows some abnormal behavior that Version 8 disallows. This parameter is deprecated and is retained for backward compatibility only.
remote_os_authent	REMOTE_OS_AUTHENT specifies whether remote clients will be authenticated with the value of the OS_AUTHENT_PREFIX parameter. This parameter is deprecated and is retained for backward compatibility only.
resource_manager_cpu_allocation	Sets the number of CPUs that Resource Manager should utilize. Default is 0. This parameter was introduced in 11.1.0.6 and deprecated in 11.1.0.7.
standby_archive_dest	STANDBY_ARCHIVE_DEST can be used to specify where archived logs received from a primary database are stored on a standby database. It is no longer necessary to set this parameter, because an appropriate location is automatically chosen.
transaction_lag	Attribute of the CQ_NOTIFICATION\$_REG_INFO object. Can be used to specify the number of transactions/database changes, by which the client is willing to lag behind the database.
Database 11g Release 2 (11.2)	
active_instance_count	ACTIVE_INSTANCE_COUNT enables you to designate one instance in a two-instance cluster as the primary instance

	and the other instance as the secondary instance. This parameter is deprecated and is retained for backward compatibility only.
parallel_io_cap_enabled	This parameter is replaced by the PARALLEL_DEGREE_LIMIT when set to IO.

Obsolete Parameters in Database 11g

Parameter	Description
Database 11g Release 1 (11.1)	
ddl_wait_for_locks	Specifies whether DDL statements (such as ALTER TABLE ... ADD COLUMN) wait and complete instead of timing out if the statement is not able to acquire all required locks.
logmnr_max_persistent_sessions	Specifies the maximum number of persistent LogMiner mining sessions that are concurrently active when all sessions are mining redo logs generated by standalone instances.
plsql_compiler_flags	Specifies a list of flags for the PL/SQL compiler as a comma-separated list of strings.
Database 11g Release 2 (11.2)	
drs_start	Enables Oracle to determine whether or not the DRMON (Disaster Recovery Monitor) process should be started.
gc_files_to_locks	A RAC parameter that controls the mapping of pre-release 9.0.1 parallel cache management (PCM) locks to datafiles.
max_commit_propagation_delay	Used when data consistency between different RAC instances must be guaranteed and immediate i.e. if commits must be seen instantaneously on remote instances.
plsql_native_library_dir	Specifies the name of a directory where the shared objects produced by the native compiler are stored.
plsql_native_library_subdir_count	Specifies the number of subdirectories created by the database administrator in the directory specified by plsql_native_library_dir.
sql_version	SQL language version parameter for compatibility issues.

DISCLAIMER: *This list was built for my personal use and I do not guarantee that it is error free. I used almost exclusively the official Oracle documentation and some Metalink notes, where the official docs were not clear enough. The latest version of this document is available at [my website](#).*