

Oracle 9i: New Features For Administrators

1Z0-030

**Demo Version
From
ITCertKeys.com
To
CertsBraindumps.com**

This study guide consists of 142(140+2) number of questions.

Question 1.

For the best performance, what should be done with blocks transferred using the Cache Fusion algorithm in a Real Application Clusters database?

- A. They should be transferred across a high-speed cluster interconnect.
- B. They should be written to disk by the sending instance and read from disk by the receiving instance.
- C. They should be granted, by the Global Cache Service, an exclusive mode resource for each instance involved.
- D. They should be transferred through the cache area of the Global Resource Director on the instance where the block is mastered.

Answer: C

Explanation:

Oracle 9i real application clusters Cache fusion introduced more robust and accurate concepts to manage collaboration between multiple instances. The Global Cache Service now handles the management functions that had been handled by the DML in the early versions of the Oracle cluster software releases. It maintains block modes for blocks in the global role and is responsible for block transfers among instances. In 9i the Global Cache element is an Oracle-Specific data structure representing a Cache fusion Resource. The Global Cache Service (GCS) handles the coordination of any changed blocks. By default, each data block in an instance's buffer cache is protected by the Global Cache Service

Question 2.

Consider this RMAN command:

```
RMAN> CONFIGURE RETENTION POLICY  
P> TO RECOVERY WINDOW OF 7 DAYS;
```

What does the command accomplish?

- A. Established the fixed number of backups that must be kept for media recovery.
- B. Establishes the fixed number of backups required to perform media recovery with 7 days.
- C. Sets the fixed number of days that backups are held for media recovery.
- D. Ensures that sufficient backups and logs will be kept so that a point-in-time recovery to any time in the last 7 days is possible

Answer: C

Explanation:

To ensure that backups will be kept for at least seven days for media recovery, you should use this CONFIGURE command:

```
CONFIGURE RETENTION POLICY TO RECOVERY WINDOW OF 7 DAYS;
```

The CONFIGURE RETENTION POLICY command creates a persistent backup retention policy. With a backup retention policy, RMAN considers backups and copies of data files and control files as obsolete based on the options specified in the CONFIGURE RETENTION POLICY command. When backups become obsolete, you can view them using the REPORT OBSOLETE command or delete them using the DELETE OBSOLETE command. This is not the same as a backup expiration. RMAN only flags a backup as expired when it performs a crosscheck and determines that a file is missing or not found. Obsolete files may remain intact until you issue the DELETE

OBSOLETE command. With the CONFIGURE RETENTION POLICY command, you have two methods for specifying the retention policy, either by stating a recovery window using the TO RECOVERY WINDOW keywords or by stating how many redundant backups you wish to keep using the TO REDUNDANCY keywords. These options are mutually exclusive. The RECOVERY WINDOW parameter of the CONFIGURE command specifies the number of days between the current time and the earliest point of recoverability. Backups do not become obsolete until they fall outside of the recovery window. Recovery Window establishes a period of time within which a point-in-time recovery must be possible.

Question 3.

To increase the availability of an Advanced Replication environment, you can add a new master to a replication group without quiescing the group. This is achieved_____

- A. Only when the new master site already has existing replication groups.
- B. By issuing a single command to create, populate, and activate the new master.
- C. When the master definition site is not the same for all of the master groups.
- D. Without impacting end users who are executing data manipulation language (DML) commands on the replicated tables.

Answer: D

Explanation:

You can either add new master sites to a master group that is running normally or to a master group that is quiesced. If the master group is not quiesced, then users can perform data manipulation language (DML) operations on the data while the new master sites are being added. However, more administrative actions are required when adding new master sites if the master group is not quiesced.

Question 4.

Examine this statement, which creates a Cartesian product of the COUNTRIES and REGION tables:

```
SQL> SELECT c.country_name,  
2> r.region_name  
3> FROM countries c, region r;
```

Oracle 9i SQL: 1999 syntax supports the same functionality with which join type?

- A. Equijoin
- B. Cross join
- C. Merge join
- D. Natural join

Answer: B

Explanation:

In Oracle 9i Cross-join is used to create an intentional Cartesian product by joining all rows in one table to all rows in another table.

Question 5.

Which three table transformations can be done by using online redefinition of tables? (Choose three).

- A. delete rows
- B. rename a column
- C. change a heap table to an index-organized table

- D. change the data type of a column from LONG to BLOB
- E. change the data type of a column from DATE to CHAR

Answer: B, C & D

Explanation:

Online table redefinition enables you to: Move the table to a different table space in the same schema .Add or drop partitioning support. Re-create the table to reduce fragmentation. Change the organization of a normal table (heap organized) to an index-organized table and vice versa Add or drop a column and to change the datatype from LONG to BLOB.

Question 5(Similar).

Which three table transformations can be done using online redefinition of tables? (Choose three.)

- A. delete rows
- B. drop a column
- C. change a heap table to an index-organized table
- D. change the data type of a column from LONG to BLOB
- E. change a range-partitioned table to a list-partitioned table.

Answer: B, C & D

Explanation:

Online table redefinition enables you to Move the table to a different table space in the same schema .Add or drop partitioning support. Re-create the table to reduce fragmentation. Change the organization of a normal table (heap organized) to an index-organized table and vice versa also add or drop a column and to change the datatype from LONG to BLOB.

Question 6.

What must you do to enable Automated SQL Execution Memory Management?

- A. Execute the DBMS_STATS.GATHER_SYSTEM_STATS procedure with appropriate value.
- B. Set the instance parameters PGA_AGGREGATE_TARGET and WORKAREA_SIZE POLICY to appropriate values.
- C. Set the instance parameter SORT_AREA_SIZE, BITMAP_MERGE_AREA_SIZE, CREATE_BITMAP_AREA_SIZE, HASH_AREA_SIZE, SORT_AREA_SIZE, and SORT_AREA_RETAINED_SIZE all to AUTO.
- D. Unset the instance parameters SORT_AREA_SIZE, BITMAP_MERGE_AREA_SIZE, CREATE_BITMAP_AREA_SIZE, HASH_AREA_SIZE, SORT_AREA_SIZE, and SORT_AREA_RETAINED SIZE.

Answer: B

Explanation:

Oracle 9i offers the option of automated PGA space management. Two new parameters have been introduced to allow the DBA to have the PGA configured dynamically by the RDBMS software removing the responsibility from the DBA. These two parameters are:

1. pga_aggregate_target
2. workarea_size_policy

PGA_AGGREGATE_TARGET or the WORKAREA_SIZE_POLICY initialization parameter must be specified is incorrect because both parameters must be set. However, the default of the WORKAREA_SIZE_POLICY parameter depends on the setting of PGA_AGGREGATE_TARGET. If PGA_AGGREGATE_TARGET contains a value, then

WORKAREA_SIZE_POLICY defaults to AUTO. So, specifying only the PGA_AGGREGATE_TARGET would be successful, but if you attempt to set the WORKAREA_SIZE_POLICY prior to setting PGA_AGGREGATE_TARGET, you will receive an "ORA-04032: pga_aggregate_target must be set before switching to auto mode" error.

Question 7.

When querying the new V\$SPPARAMETER view, the column ISSPECIFIED displays the value TRUE for a number of a parameter. What does this means?

- A. The value for the parameters is TRUE
- B. The parameters are specified in the PFILE parameter file.
- C. The parameters are specified in the SPFILE parameter file.
- D. Changes to the parameters are only in the currently running instance.

Answer: C

Explanation:

V\$SPPARAMETER purpose is to allow the DBA to view the setting of the parameters in the instance SPFILE. The ISSPECIFIED column has a value of TRUE if the parameter was specified in the SPFILE.

Question 8.

Examine the following parameter settings from an initialization (init.ora) file:

```
DB_CREATE_FILE_DEST          =      '/u01/oradata'
DB_CREATE_ONLINE_LOG_DEST_1  =      '/u02/oradata'
DB_CREATE_ONLIN_LOG_DEST_2   =      '/u03/oradata'
```

If you create an Oracle Managed Files (OMF) database using these settings, what is the results?

- A. The data files, temp files, and control file will be on device /u01; one online redo log group will be on device /u02; another redo log group will be on device /u03.
- B. The data files and temp files will be on device /u01; one copy of the control file and one online redo log group will be on device /u02; another copy of the control file and a second redo log group will be on device /u03.
- C. The data files, temp files, online redo log files, and control file will be on device /u01; multiplexed copies of the archive log files will be created; one set on device /u02, and another set on device /u03.
- D. The data files and temp files will be on device /u01; one copy of the control file and the first member in each online redo log group will be on device /u02; another copy of the control file and a second member of each redo log group will be on device /u03.

Answer: D

Explanation:

The **db_create_file_dest** parameter defines the file system where OMF and Temp files are to be located. where as **db_create_online_log_dest_n** parameter defines the file system location where Oracle managed online redo logs files and control files are to be created. The **n** Value is replaced by a number 1-5.This allows for up to 5 multiplexed copies of each redo log group member, and up to 5 copies of the control files to be created on separate derives.

Question 9.

Examine the statement:

```
ALTER TABLE sales
EXCHANGE PARTITION q1_2000
WITH TABLE sales_q1_2000
UPDATE GLOBAL INDEXES
```

PARALLEL (DEGREE 4);
What is the result of the statement?

- A. All valid global indexes on the SALES table will retain their USABLE status after the statement completes.
- B. Nonpartitioned indexes on the SALES table will be marked UNUSABLE, while valid partitioned global and local indexes will retain their USABLE status.
- C. Valid indexes on the exchanged table SALES_Q1_2000 will remain in the USABLE state, while all indexes associated with the new partition Q1_2000, will be marked UNUSABLE, including any global indexes on SALES.
- D. Global indexes on the SALES table will be maintained concurrently with the exchange operation if the statement executes with the desired degree of parallelism (four) otherwise it will execute serially and mark the global indexes UNUSABLE.

Answer: A

Explanation:

By default, many table maintenance operations on partitioned tables invalidate i.e. mark UNUSABLE global indexes. You must then rebuild the entire global index or, if partitioned, all of its partitions. Oracle enables you to override this default behavior if you specify UPDATE GLOBAL INDEXES in your ALTER TABLE statement for the maintenance operation it tells Oracle to update the global index at the time it executes the maintenance operation DDL statement. The entire global index is updated in conjunction with the base table operation. You are not required to later and independently rebuild the global index.

Question 10.

You use the TABLESPACE option of the Export command to export, not transport, the contents of a set of tablespaces. What are the only contents of the dump file?

- A. 1. non-partitioned tables in the tablespace set
2. partitioned tables that have all their partitions in the tablespace set
3. indexes on all the exported tables
- B. 1. non-partitioned tables in the tablespace set
2. partitioned tables that have at least one partition in the tablespace set
3. indexes on all the exported tables
- C. 1. non-partitioned tables in the tablespace set
2. partitioned tables that have all their partitions in the tablespace set
3. indexes on the exported tables if the indexes are in the tablespace set or have at least one partition in the tablespace set.
- D. 1. non-partitioned tables in the tablespace set
2. partitioned tables that have at least one partition in the tablespace set
3. tables (including all partitions, if they exist) that have an index in the tablespace set
4. indexes on all the exported tables

Answer: B

Explanation:

In oracle 9i when a TABLESPACE parameter is used with the export utility then all tables included in the tablespace including table those have one partition in the tablespace are exported and all indexes on exported whether they are in that tablespace or in any other tablespace will be exported to the dump file. Non-partitioned tables and partitioned tables and all indexes in the table can be exported through dump file.

Question 11.

The new time and date data type support greater precision for capturing fractional seconds and time zone information. This gives the Oracle9i database greater flexibility in supporting locally preferences.

What is the default degree of precision of fractional seconds when specifying the TIMESTAMP data type?

- A. 6 digits
- B. 9 digits
- C. 12 digits
- D. 22 digits

Answer: A

Explanation:

All values of TIMESTAMP as well as time zone displacement value, where **fractional_seconds_precision** is the number of digits in the fractional part of the SECOND datetime field. Accepted values are 0 to 9. The default is 6.

Question 12.

Before adding a nonstandard block size to your database, which preparatory task must you complete?

- A. Define a new buffer pool by setting the DB_nK_CACHE_SIZE parameter where n is the new block size.
- B. Redefine the buffer pool by setting the DB_mK_CACHE_SIZE and DB_nK_CACHE_SIZE parameter, where m is the standard block size and n is the new block size.
- C. Rebuild the control file using an instance with an initialization file containing two DB_BLOCK_SIZE parameter, one for the standard and one for the new block size.
- D. Rebuild the control file using an instance with an initialization file containing two DB_nK_CACHE_SIZE parameters, one for the standard and one for the new block size.

Answer: A

Explanation:

The DB_CACHE_SIZE parameter specifies the size of the default cache for the database's standard block size. To create and use tablespaces with block sizes different than the database's standard block sizes you must configure a separate cache for each block size used. The DB_nk_CACHE_SIZE parameter can be used for this purpose.

Question 13.

Which two procedures or functions are part for the DBMS_METADATA package? (Choose two)

- A. GET_DDL
- B. GET_XML
- C. GET_TYPE
- D. GET_VIEW
- E. GET_TABLE

Answer: A & B

Explanation:

The Metadata Application Programming Interface (API) would be used to retrieve DDL statements from the data dictionary in XML format. Oracle9i provides the DBMS_METADATA package that allows users to retrieve information from the data dictionary.

The **dbms_metadata** package allows several functions. Principal ones for nonprogrammatic use are **get_ddl** and **get_xml**. The **get_ddl** function returns the object information as SQL-runnable DDL, where as the **get_xml** returns the data as object metadata in XML format.

Question 14.

Which two are true regarding the Automatic Undo Management feature? (Choose two)

- A. PMON is responsible for shrinking undo segments as needed.
- B. SMON is responsible for shrinking undo segments as needed.
- C. To use this feature, you must create undo segments into an undo tablespace.
- D. Oracle9i automatically creates a fixed number of undo segments per undo tablespace.
- E. In a Real Application Cluster environment, you must create one undo tablespace per opened instance.

Answer: B & D

Explanation:

Oracle9i introduces a new feature called Automatic Undo Management (AUM). This feature allows rollback data to be managed more easily. Instead of managing space for each rollback segment, an undo tablespace can be created with enough disk space for the entire instance. Oracle can create new undo segments or shrink existing ones as needed. When automatic undo management is being used, you cannot CREATE, DROP, or ALTER undo segments. In a Real Application Clusters environment, each instance must have its own undo tablespace.

When using Oracle9i's Automatic Undo Management feature, SMON automatically shrinks undo segments as needed. SMON shrinks the undo segments every twelve (12) hours to eliminate space from inactive undo segments. SMON also shrinks the undo segments each time a process needs to get additional space from another undo segment.

Question 15.

A switchover to a standby database differs from a failover to the standby in a number of ways.

What are two of the key differences that characterize a switchover? (Choose two)

- A. The operation is planned and causes no data loss.
- B. The standby database must be running in a no data divergence mode.
- C. The old primary database can take on the standby role and be available for a future switchover.
- D. The primary and standby database must be using exactly the same release and patch level of Oracle 9i.
- E. Either the primary database's online or archived log files, but not necessarily both, must be available.

Answer: A & C

Explanation:

A database switchover is not the same as a database failover. A failover usually occurs as a result of an unplanned outage of the primary database. A switchover, on the other hand, is a planned operation. For failover, the primary database role is transferred to a standby database. The original primary database cannot be used as a standby, and a new standby must be created.

Question 16.

Which command can you use to alter the active size of the SGA dynamically, where the entered value differs from the current one?

- A. ALTER SYSTEM SET SGA_MAX_SIZE=500M;

- B. ALTER SYSTEM DROP DB_2K_CACHE_SIZE;
- C. ALTER SYSTEM SET DB_CACHE_SIZE=AUTO;
- D. ALTER SYSTEM SET SHARED_POOL_SIZE=40M;

Answer: D

Explanation:

Oracle9i has made more instance parameters dynamic, including DB_CACHE_SIZE and SHARED_POOL_SIZE:

```
ALTER SYSTEM SET SHARED_POOL_SIZE = 10000000;  
ALTER SYSTEM SET DB_CACHE_SIZE = 10000000;
```

Question 17.

When adding a new master site in an Advanced Replication environment without quiescing replication activity, you can instantiate the database at the new master site by using one of two methods.

Which two methods can be used to perform this step? (Choose two)

- A. Perform change-based recovery on the new master site.
- B. Perform a cancel-based recovery on the new master site.
- C. Perform a time-based recovery on the new master site.
- D. Transport the relevant tablespace set from the master definition site to the new master site.
- E. Perform an export using the FLASHBACK_SCN option on the master definition site and then import the dump file to the new master database.

Answer: A & E

Explanation:

With Oracle9i's Advanced Replication enhancements, you can now add new master sites in an Advanced Replication environment without quiescing. One of two methods can be used to accomplish this task. You can add new masters and then perform change-based recovery on the new master site, or you can add the new masters and then export and import. If the latter option is chosen then you should include the FLASHBACK_SCN parameter on the export.

Question 18.

Which operation causes an index to be considered "used" while monitoring is turned on for that particular index?

- A. Only when the index is specified in a hint.
- B. When data is fetched by an execution plan that uses that index.
- C. When a statement is executed and the execution plan contains a reference to the index.
- D. When a statement is parsed and the resulting execution plan contains a reference to the index.

Answer: D

Explanation:

The USED column of the V\$OBJECT_USAGE view indicates whether a SQL statement has been parsed that used the index in its execution plan. The USED column is only updated at the time a statement parses. Therefore, if you issue an ALTER INDEX statement with the MONITORING USAGE clause and the statement has already begun executing, the USED column will not be updated even though the index is being accessed by the running statement.

Question 19.

What does the Character Set Scanner command line utility do?

- A. reports only on any Unicode character conversions required
- B. performs character set conversions to the new specified character set
- C. scans on any character set conversion required, then converts the data
- D. checks for any character definitions that will fail conversion to the new character set.

Answer: D

Explanation:

The Character Set Scanner is a separate utility program available in Oracle9i. The program is used to determine if changes to the character set will result in any errors. When converting to a different character set, data can be truncated or corrupted or a character set mismatch may occur. Using the Character Set Scanner utility allows you to review these potential problems and make corrections prior to initiating the conversion.

Question 20.

What is true regarding fact tables having bitmap join indexes defined on them?

- A. You can create a bitmap join index on a temporary table.
- B. You can create a bitmap join index on an index-organized table.
- C. Tables participating in the join condition of the bitmap join index definition cannot be updated concurrently.
- D. You cannot have more than one dimension table being part of the join index definition.

Answer: C

Explanation:

Oracle9i introduces a new type of index called a Bitmap Join Index. A Bitmap Join Index is a bitmap index defined for the join of two or more tables. Bitmap Join indexes would be most useful in a data warehousing environment where a fact table and dimension tables are often queried in a single statement. Using a Bitmap Join index provides for better query performance in this type situation. To define a Bitmap join index, the BITMAP keyword is included in the CREATE BITMAP INDEX statement and a FROM and WHERE clause define the join condition. When a Bitmap Join index is defined, only one table referenced by the index can be updated concurrently and no table can appear twice in the join condition.

Question 21.

Which statement describes the distinction between Fine-Grained Audit (FGA) and Fine-Grained Access Control (FGAC)?

- A. FGA is an extension of FGAC.
- B. FGAC prohibits access by row; FGA enables access by row.
- C. Both use Application Context, but only FGAC can use the Secure Application role.
- D. FGA tracks when sensitive rows have been accessed; FGAC prevents access to sensitive rows.

Answer: D

Explanation:

Fine grained Auditing is the new feature of 9i Enterprise Edition that allows you to audit **Select** statements.

Fine-Grained Access Control, not Fine-Grained Auditing prevents access to specific rows of data. With Fine-Grained Auditing, only SQL access is audited

Question 22.

You need to version enable the MORTGAGE table containing customer mortgage trend information. Which syntax version enables the MORTGAGE table?

- A. DBMS_LT.ENABLETABLE ('MORTGAGE');
- B. DBMS_WM.ENABLE_TABLE ('MORTGAGE');
- C. DBMS_LT.ENABLEVERSIONING ('MORTGAGE');
- D. DBMS_WM.ENABLEVERSIONING ('MORTGAGE');

Answer: D

Explanation:

Every time a change is made to a version-enabled table a new row is created to reflect the change, with all new rows being stored in the same tablespace. Since only changed rows are stored, the storage overhead is minimized. The topmost workspace is always used to create a new workspace, irrespective of where it is in the workspace hierarchy. A table can be version-enabled using:

```
EXEC DBMS_WM.ENABLEVERSIONING('MORTGAGE');
```

This will create a view for the version enable table which will be used instead of triggers to perform all operation against version enable table.

Question 23.

Which three resource plan parameters are used to enable Automatic Consumer Group Switching to estimate job execution times automatically, and to perform a switch to a low priority consumer group before they are initiated? (Choose three).

- A. SWITCH_TIME
- B. SWITCH_GROUP
- C. MAX_SESS_POOL
- D. SWITCH_ESTIMATE

Answer: A, B & D

Explanation:

Automatic Consumer Group Switching: Oracle8i allowed resource intensive transactions to be manually switched between consumer groups to give them less priority. Oracle9i improves on this to allow the session to be automatically switched when it exceeds a specified threshold:

- **SWITCH_GROUP** - Defines the consumer group to switch to if the threshold is exceeded.
- **SWITCH_TIME** - Defines the threshold time.
- **SWITCH_ESTIMATE**- Defines the estimated threshold time.

Question 23(Similar).

Automatic Consumer Group Switching is an important new feature of the Database Resource Manager.

Which three plan directive parameters are used to control this feature? (Choose three.)

- A. SWITCH_TIME
- B. SWITCH_GROUP
- D. MAX_SESS_POOL
- D. SWITCH_ESTIMATE

Answer: A, B & D

Explanation:

Using automatic consumer group switching feature of Oracle 9i, DBA can specify a criteria which will cause the database resource manager to automatically switch the consumer group of long sessions and automatic consumer switching is controlled by three features as:

Switch group which specifies the consumer group to which this session is switched if the other criteria are met these criteria are other features and then switch time specifies the length of time that a session can execute before it is switched to another consumer group and third feature is switch estimate which tells how much time a session will go on execute before it is switched to another consumer group.

Question 24.

What are the two methods used to add a second, or subsequent block size to an existing database? (Choose two.)

- A. Create a new tablespace and include a block size specification.
- B. Transport a tablespace with a different block size into the database.
- C. Re-create the control file to specify a new block size for specific data files.
- D. Take a tablespace offline, and then put it back online with a new block size specification.

Answer: A & B

Explanation:

You can create tablespaces of different block sizes than the standard database block size specified by the DB_BLOCK_SIZE initialization parameter. However, your buffer cache in SGA memory must be configured for the nonstandard block sizes.

You can specify a non standard block size when creating a new tablespace. This new tablespace can be created with different tablespace.

The ability to specify multiple block sizes for your database is especially useful if you are transporting tablespaces between databases. You can, for example, transport a tablespace that uses a 4K block size from an OLTP environment to a datawarehouse environment that uses a standard block size of 8K. If you are transporting a tablespace of a different block size than the standard block size of the database receiving the tablespace set, then you must first have a DB_nK_CACHE_SIZE initialization parameter entry in the receiving database's parameter file.

Question 25.

What is required for the Character Set Scanner to operate?

- A. Oracle Enterprise Manager
- B. A special schema installed with csminst.sql
- C. A separate tablespace called cs_scan to store the results.
- D. A GUI capable workstation, because it uses a Java interface.

Answer: B

Explanation:

Before using the Character Set Scanner, you must execute the csminst.sql script. This script will create a schema that is used by the Character Set Scanner. Once this has been done, you can run the Character Set Scanner, review the three output files generated by the Character Set Scanner, and then perform any necessary conversions.

Question 26.

Examine the statement

```
SQL> CREATE TABLESPACE user_DATA
2> EXTENT MANAGEMENT LOCAL
3> SEGMENT SPACE MANAGEMENT AUTO;
```

Which two assumptions must be true for this statement to execute successfully? (Choose two.)

- A. Oracle Managed Files are used for this instance.
- B. The USER_DATA tablespace is managed using FET\$/UET\$ tables.
- C. The COMPATIBLE initialization parameter must be 9.0.0 or higher.
- D. Space within segments in the USER_DATA tablespace is managed with freelists.

Answer: A & C

Explanation:

Without oracle managed files OMF segment space management 'Auto' is impossible which is only possible in 9i. Automatic segment-space management is specified at the tablespace level using the SEGMENT SPACE MANAGEMENT AUTO option.

Question 27.

The Oracle9i LogMiner needs a data dictionary copy to display the object names.

From which three locations can LogMiner retrieve the data dictionary information? (Choose three).

- A. separate supplemental log file
- B. information in the control file
- C. the current online data dictionary
- D. a Data Dictionary copy stored in the redo log file.
- E. a Data Dictionary copy stored in an operating system flat file
- F. redo information stored within the redo vectors of the DML operation

Answer: C, D & E

Explanation:

Oracle9i's LogMiner utility allows you to query redo log files using SQL. However, to fully translate the log data, LogMiner requires a source dictionary. Prior to Oracle9i, you could only use a dictionary that had been extracted to a flat file. With Oracle9i, LogMiner can use a source dictionary that has been extracted to an operating system file, a dictionary that has been extracted to the redo logs, or the current online data dictionary. If no dictionary is specified, then LogMiner shows redo log information in hexadecimal format.

You cannot extract a data dictionary to a supplemental log file. You can use the ADD SUPPLEMENTAL LOG DATA clause to include additional columns of data in the log stream for LogMiner to use, but the data dictionary is not accessed from a supplemental log file

Question 28.

What type of protection is implemented when the DBA issues the ALTER DATABASE SET STANDBY DATABASE PROTECTED; command?

- A. The standby database is protected against write operations.
- B. The primary database is protected against write operations.
- C. The primary database is protected against data loss and data divergence.
- D. The standby database is protected against data loss and data divergence.

Answer: C

Explanation:

Protect Primary Database: Now that Data Guard is configured and running the primary database can be prevented from applying updates unless the update has been sent to at least one standby location. Connect to the primary database and execute:

```
ALTER DATABASE SET STANDBY DATABASE PROTECTED;
```

When in PROTECTED mode, online redo logs, rather than archived redo logs, are sent to the standby. The standby cannot diverge at all from the primary database. When a commit occurs on the primary, control is not returned to the user until all the data is available at the standby no data loss will occur. The standby database cannot be made read-only. Doing so would make the standby unavailable, and the primary database would shut down.

Question 29.

In a Data Guard Switchover operation, which command must you execute on the database that is being switched from the standby to the primary role?

- A. ALTER DATABASE MOUNT STANDBY DATABASE;
- B. ALTER DATABASE BACKUP CONTROLFILE TO TRACE;
- C. ALTER DATABASE COMMIT TO SWITCHOVER TO PHYSICAL PRIMARY;
- D. ALTER DATABASE COMMIT TO SWITCHOVER TO PHYSICAL STANDBY;

Answer: B

Explanation:

Database Switchover: A database can be in one of two mutually exclusive modes (primary or standby). These roles can be altered at runtime without loss of data or resetting of redo logs. This process is known as a Switchover and can be performed using the following statements:

```
ALTER DATABASE COMMIT TO SWITCHOVER TO PHYSICAL PRIMARY;
```

This statement switches the standby database to the primary role. Oracle ensures that all log files have been received, closes the database if necessary. The selected standby database will now be converted into a primary database. The current standby control file will not be backed up to the current trace file. You can always reconstruct the standby control file from the original primary database, which is a new standby database that has a copy of control file and to take the backup of control file use the command.

```
ALTER DATABASE BACKUP CONTROLFILE TO TRACE
```

Question 30.

In order to speed foreign key creation, Oracle9i will cache the first _____ primary key values only where there are multirow DML statements.

- A. 32
- B. 128
- C. 256
- D. 512

Answer: C

Explanation:

Oracle9i provides for caching of primary keys. When inserting records with foreign keys, the matching primary key must be looked up. Oracle9i stores the first 256 primary key values in memory. Therefore, foreign key creation is faster. To reduce overhead, caching occurs when the

second row is processed. This prevents the need to cache primary key values for single-row DML statements. The first 256 primary keys are cached for multi-row DML statements.

Question 31.

Examine this fragment from a SQL*Plus session:

```
SQL> SELECT name, value FROM v$sysstat  
2> WHERE name LIKE `work area executions%`;
```

NAME	VALUE
work area executions – optimal	1544
work area executions – onepass	11
work area executions – multipass	1038

What can you conclude about the setting of the PGA_AGGREGATE_TARGET initialization parameter?

- A. You cannot conclude anything because these statistics are not related to the PGA_AGGREGATE_TARGET parameter.
- B. It may be set too low because there are many more multipass work area executions than one pass executions.
- C. It may be set too high because there are many more multipass work area executions than one pass executions.
- D. It is probably at its optimal setting because the ratio of one pass to multipass work area executions is below 0.10 (ten percent).

Answer: B

Explanation:

If the number of multi-pass executions is non zero, then consider increasing the value of PGA_AGGREGATE_TARGET parameter.

When the work area executions - multipass statistic contains a high value, then you should consider increasing the PGA_AGGREGATE_TARGET value. The work area executions - multipass statistics indicates the total number of work areas running in more than one pass. This means that the amount of available PGA was exceeded. Increasing the PGA_AGGREGATE_TARGET value will allow for the system to be tuned more efficiently.

Question 32.

The enterprise data warehouse routinely receives data from various production OLTP databases. The database block size on the data warehouse is 8192, and in the OLTP databases it may be 2048, 4096 or 8192. The current SGA is 112 MB.

Oracle9i has the ability to transport tablespaces with difference block sizes.

Which two parameters are required in order to “plug in” the OLTP tablespaces into the warehouse database? (Choose two).

- A. DB_2K_CACHE_SIZE
- B. DB_4K_CACHE_SIZE
- C. DB_8K_CACHE_SIZE
- D. DB_2K_BLOCK_BUFFERS
- E. DB_4K_BLOCK_BUFFERS
- F. DB_8K_BLOCK_BUFFERS

Answer: A & B

Explanation:

Oracle 9i supports tablespaces with different block sizes. This allows large tables and indexes to have a larger block size than smaller objects. This is especially useful in hybrid databases where DSS transactions benefit from large block sizes, whilst OLTP operations are best suited to smaller block sizes. It also means that tablespaces can be transported between databases with differing block sizes. The database is created with a standard block size and up to 5 non-standard block sizes.

An individual buffer cache must be defined for each non-standard block size used
You cannot define a non-standard block size to be the same size as the standard block size

Question 33.

Which two are true regarding a list-partitioned table? (Choose two)

- A. It is possible to partition an index-organized table using the LIST method.
- B. It is NOT possible to have multiple columns as the partition key of a list-partitioned table.
- C. The keyword NULL can be specified as a partition literal value for one partition of a list-partitioned table.
- D. The keyword MAXVALUE can be specified as a partition literal value for one partition of a list-partitioned table.

Answer: B & C

Explanation:

Use list partitioning when you require explicit control over how rows map to partitions. You can specify a list of discrete values for the partitioning column in the description for each partition. Only one column in the table can be defined in the partition by list clause and Key word NULL can be specified as a partition value.

Question 34.

What are two key features of Automatic Undo Management? (Choose two)

- A. Oracle manages all undo space automatically.
- B. The SYSTEM tablespace is never used to store automatic undo segments.
- C. Undo tablespace contain all undo information not stored in the SYSTEM rollback segment.
- D. Rollback segments are gradually replaced with automatic undo segments as the instance gathers statistics about rollback requirements.
- E. The number of automatic undo segments available to an instance is determined by the number of undo tablespaces defined in its initialization parameter file.

Answer: A & C

Explanation:

Oracle9i introduces a new feature called Automatic Undo Management (AUM). This feature allows rollback data to be managed more easily. Instead of managing space for each rollback segment, an undo tablespace can be created with enough disk space for the entire instance. Then, transactions are able to use the undo tablespace dynamically as needed. Oracle can create new undo segments or shrink existing ones as needed. When automatic undo management is being used, you cannot CREATE, DROP, or ALTER undo segments.

Question 35.

Oracle Enterprise Manage (OEM) can generate database reports. Which statement about this feature is true?

- A. Generating reports requires a repository.

- B. Report generation is scheduled in the job system.
- C. The report is placed in a table for suitable selects.
- D. The reports are in HTML format for viewing in a browser.

Answer: D

Explanation:

When a report is generated, the output is stored in a file, not in a table. Reports can be generated in HTML, text, or Comma-Separated Values format.

Question 36.

Which four are true regarding the Workspace Manager in Oracle9i? (Choose four)

- A. Automatically versions all tables
- B. Automatically installed with Oracle9i.
- C. Merges changes with parent rows or discards changes
- D. Provides mechanism to identify and resolve conflicts.
- E. Allows for version enabling tables by use of a packaged procedure

Answer: B, C, D & E

Explanation:

Workspace Manager is installed in Oracle 9i by default. Workspaces allow multiple transactionally consistent environments to exist within one database. This allows several departments or functional areas to work against a single schema without interfering with data from other groups. Changes to version-enabled tables are captured as new rows within the workspace. These changes are invisible to other workspaces until they are merged into a parent workspace. Multiple versions of a row can exist within a workspace, with the active or current row being the one to which changes are currently being made.

In a workspace hierarchy consisting of **Live->PreProduction->Development** workspaces, the **Development** workspace can see all row changes made in the **PreProduction** workspace, along with all committed data from non-version-enabled tables belonging to the **Live** workspace. In addition it can see data from version-enabled tables in **Live** as they were when the **PreProduction** workspace was created. Once a workspace is refreshed, all changes can be cascaded down the hierarchy.

Conflicts are detected automatically before changes are merged into the parent workspace. These conflicts can be corrected using Enterprise Manager or a programmatic API.

Question 37.

What statement is true about taking a database default temporary tablespace offline?

- A. You cannot take default temporary tablespace offline.
- B. The database must be quiesced before you perform this action.
- C. You must ensure that all users are assigned to an alternative temporary tablespace first.
- D. You must first ensure that the default temporary tablespace has the permanent tablespace characteristic.

Answer: A

Explanation:

Oracle9i introduces the new default temporary tablespace feature. This allows you to specify a database-wide default temporary tablespace either when creating the database or using the ALTER DATABASE statement. If users are added and not assigned a temporary tablespace, the

tablespace specified as the default temporary tablespace is assigned to the user. Once defined, the default temporary tablespace CANNOT be taken offline.

You cannot make the default temporary tablespace permanent or take it offline but you can drop the default temporary tablespace in that situation SYSTEM tablespace will be used as default temporary tablespace

Question 38.

User defined events in Oracle Enterprise Manager return status, and possibly values, to the Oracle Enterprise Manager Event Monitor. What are user defined events?

- A. Database based scripts (PL/SQL, Java) run by the Agent.
- B. Database Event Triggers, which issues Alerts to the Agent.
- C. Scripts run by the Agent in any language on the host server
- D. User written PL/SQL programs in the database, which the OEM console can execute and display the results.

Answer: C

Explanation:

User Defined Events allow you to integrate any monitoring script with the Event System. The Agent will run any specified script which you may define to check for conditions specific to your environment. The script itself can be written in any language, as long as the runtime requirements needed by the Agent to run the script are available on the monitored node

Question 39.

Which two statements are true regarding the use of UTF-16 encoding? (Choose two.)

- A. Enables easier loading of multinational data.
- B. Uses a fixed-width multibyte encoding sequence
- C. Asian characters are represented in three characters
- D. Uses a variable-width multibyte encoding sequence
- E. European characters are represented in one or two bytes.

Answer: A & B

Explanation:

UTF-16 Encoding: This is the 16-bit encoding of Unicode. It is a 2 byte fixed-width encoding in which the character codes 0x0000 through 0x007F have the same meaning as ASCII.

One Unicode character is 2-bytes in this encoding. Characters from all scripts are represented in 2 bytes.

By default data store in UTF-16 encoding in SQL NCHAR datatype which store data in fixed width and also enable easier loading of multinational data.

Question 40.

You want to drop the TBS1 tablespace from your database. You also want to delete the corresponding data files automatically, and not have to do it manually. What should you do?

- A. Use the DROP DATAFILE command.
- B. Use the DROP TABLESPACE command.
- C. Ensure that all database files are Oracle Managed Files before using the DROP TABLESPACE command.
- D. Ensure that the DB_FILE_CREATE_DEST initialization parameter is set before using the DROP TABLESPACE command.

Answer: B

Explanation:

In Oracle 9i nothing is required for dropping the tablespace except a simple command of DROP TABLESPACE and Oracle will handle the rest of the operation.

Question 41.

Oracle 9i provides a database package called dbms_redefinition to perform an online rebuild of a table. Which two steps are you recommended to do prior to issuing the dbms_redefinition.start_redef_table procedure call? (Choose two)

- A. Grant privileges on the interim table.
- B. Invoke the dbms_redefinition.can_redef_table procedure.
- C. Invoke the dbms_redefinition.sync_interim_table procedure.
- D. Create any triggers, indexes, or constraints on the interim table
- E. Create an empty interim table with all the desired characteristics

Answer: B & E

Explanation:

In highly available systems, it is occasionally necessary to redefine large "hot" tables to improve the performance of queries or DML performed against these tables. Oracle provide a mechanism to redefine tables online. This mechanism provides a significant increase in availability compared to traditional methods of redefining tables that require tables to be taken offline
Create an empty interim table and invoke the DBMS_REDEFINITION.can_redef_table procedure.

Question 42.

During scheduled maintenance for a database that runs 24 hours a day, 7 days a week, the DBA wants to bring online a large undo tablespace, smu_tbs2, while application users are performing DML. The current ONLINE undo tablespace is smu_tbs1.

The DBA issues the following command:

```
SQL> ALTER SYSTEM SET undo_tablespace = smu_tbs2;
```

What happens?

- A. Because the database is operating in Automatic Undo Management (AUM) mode, Oracle will dynamically balance a new rollback usage request across segments in both smu_tbs1 and smu_tbs2.
- B. All new transactions will use rollback segments in smu_tbs2, and any exiting sessions with active transactions will get an ORA-3113:end-of-file on communication channel error.
- C. The ALTER SYSTEM command will fail if there are sessions with active transactions, so the DBA must have the database in restricted mode to perform the operations.
- D. All undo segments in the smu_tbs1 tablespace will have v\$rollstat.status = 'PENDING OFFLINE' until all existing transactions have been committed or rolled back. All new transaction will use smu_tbs2.
- E. In order for the application users not to be impacted by the maintenance operation, the DBA will need to dynamically change undo_management = MANUAL in order for users to use rollback segment undo while the undo tablespace is modified.

Answer: D

Explanation:

You can switch undo tablespaces even if transactions are executing. The switch proceeds and any new user transactions use the new undo tablespace. Any pending transactions continue to use the old undo tablespace. The tablespace is given the status of PENDING OFFLINE until all

previous user transactions complete. While the tablespace has this status, neither the current instance nor any other instance can use it. Once the transactions have completed, the tablespace returns to an OFFLINE status and is available for use by other instances

Question 43.

You are using automatic space-managed segments, and you wish to change the value of PCTFREE for a given table. What can you do to ensure the change is made in the bitmap structure?

- A. Use the ALTER command, then execute the dbms_space package.
- B. Use the ALTER command, then execute the dbms_repair.segment_fix_status procedure.
- C. Use the ALTER command to change the PCTFREE value of automatic space-managed segments.
- D. It is not possible to change the PCTFREE value of automatic space-managed segments.

Answer: B

Explanation:

With Automatic Space Managed segments, the bitmap representing the status of the data blocks is not automatically updated when you issue an ALTER TABLE statement. After issuing the ALTER TABLE statement, you should execute the DBMS_REPAIR.SEGMENT_FIX_STATUS procedure

Question 44.

The DBMS_STATS.GATHER_SYSTEM_STATS routine collects system statistics in a user-defined time frame. Which four statements are true of the GATHER_SYSTEM_STATS routine? (Choose four)

- A. You can use the GATHERING_MODE => 'END' option to end system statistics collection.
- B. You can use the INTERVAL option to specify a fixed interval to collect statistics.
- C. You can use the GATHERING_MODE => 'START' option to enable manual statistics collection.
- D. The routine allows the optimizer to consider a system's CPU and I/O utilization and performance.
- E. The routine allows the optimizer to invalidate existing cached plans.

Answer: B, C, D & E

Explanation:

The DBMS_STATS built-in package provides procedures to assist in the collection and use of system, schema, database, and object statistics. Using the GATHER_SYSTEM_STATS procedure, system statistics can be gathered for a time period. The specification of the GATHER_SYSTEM_STATS procedure is:

```
DBMS_STATS.GATHER_SYSTEM_STATS (  
  gathering_mode VARCHAR2 DEFAULT 'INTERVAL',  
  interval INTEGER DEFAULT 60,  
  statab VARCHAR2 DEFAULT NULL,  
  statid VARCHAR2 DEFAULT NULL,  
  statown VARCHAR2 DEFAULT NULL);
```

GATHERING_MODE - The GATHERING_MODE parameter can either be INTERVAL, STOP, or START. When INTERVAL is specified (or no value is specified), you can provide an additional INTERVAL parameter in minutes. The collections will be created/updated for the specified time interval. When using the START and STOP values for the GATHERING_MODE parameter, the INTERVAL parameter is ignored. You can specify to start and stop the gathering of statistics.

STATTAB - This parameter identifies a user-defined table where statistics should be saved. If no table is specified, statistics are captured directly to the data dictionary.

STATID - This parameter is an optional identifier to be saved with the statistics. Using this parameter is helpful if you are collecting multiple sets of statistics for comparison.

STATOWN - This parameter specifies the owner of the table to which the statistics are captured, if the owner is not the current user.

In this scenario, the GATHER_SYSTEM_STATS procedure was executed only with a GATHERING_MODE value of 'START'. Because no table was specified, statistics collection begins and statistics are captured directly to the data dictionary. Therefore, the option stating that statistics are captured directly to the data dictionary is correct.

You can use GATHER_SYSTEM_STATS(GATHERING_MODE=>'STOP') to stop gathering earlier when scheduled.

Question 45.

Which three parameters are ignored when creating an automatic space management segment? (Choose three).

- A. PCTFREE
- B. PCTUSED
- C. INTRANSWER
- D. MAXTRANSWER
- E. FREELISTS
- F. FREELIST GROUPS

Answer: B, E & F

Explanation:

Automatic space management segment eliminates the need to specify the PCTUSED, FREELISTS and FREELISTS GROUPS when defining the objects in the tablespace.

Question 46.

What are three characteristics of the Data Guard Broker? (Choose three.)

- A. It runs as a process called DMON.
- B. It can be controlled through the Data Guard Manager GUI and command-line interface.
- C. It is started by setting the DRS_START parameter to TRUE in your initialization file.
- D. It uses a configuration file stored on the primary database server and, optionally, on one of your standby database servers.

Answer: A, B & C

Explanation:

When the DRS_START=TRUE parameter is specified, the DMON process is started when the instance starts. The DMON process is known as the Data Guard Broker. The Data Guard Broker allows you to configure, control, and monitor the Data Guard environment.

Users can choose one of two different interfaces to perform role changes such as failover, that is, having the standby take over production processing from the primary database. One option is to use the new Oracle Enterprise Manager Data Guard Manager. It provides a graphical user interface (GUI) for most configuration and set-up tasks, as well as operational functions. A command-line interface (CLI) is also available. It provides access to both basic monitoring and all commands needed to make role changes, as well as the ability to configure and set up a Oracle9i Data Guard environment.

Question 47.

What do you need to do in order for the Database Resource Manager to estimate, in advance, the execution time of a statement?

- A. Set the TIMED_STATISTICS initialization parameter.
- B. Gather optimizer statistics on related objects.
- C. Set the TIMED_OS_STATISTICS initialization parameter.
- D. Nothing.

Answer: B

Explanation:

Database Resource Manager can also use optimizer statistics to estimate the execution time of a statement prior to the statement's execution. When the MAX_EST_EXEC_TIME parameter is specified, a limit is imposed based on optimizer statistics.

Question 48.

You are attempting to create an Oracle-Managed Files (OMF) tablespace in a production database with the following statement and receive the following error message:

```
CREATE TABLESPACE tbs1;
```

```
ORA-02199: missing DATAFILE/TEMPFILE clause
```

```
oerr ora 2199
```

```
02199, 00000, "missing DATAFILE/TEMPFILE clause"
```

```
*Cause: A CREATE TABLESPACE statement has no DATAFILE/TEMPFILE clause.
```

```
*Action: Specify a DATAFILE/TEMPFILE clause.
```

What is the corrective action to create the OMF based tablespace?

- A. Issue the CREATE TABLESPACE tbs1 DATAFILE SIZE 10M; command.
- B. Issue the CREATE TABLESPACE tbs1 EXTENT MANAGEMENT ORACLE; command.
- C. Set the db_create_file_dest parameter using the ALTER SESSION command and re-issue the statement.
- D. An OMF tablespace is not allowed; only control files and/or redo log files can be created by OMF.

Answer: C

Explanation:

DB_CREATE_FILE_DEST sets the default location for datafile, control file, and online log creation. you can specify a file system directory as the default location for the creation of datafiles, control files, and online logs. However, the directory must already exist; Oracle does not create it. So in this above mentioned case with this command you just have to set DB_CREATE_FILE_DEST parameter by issuing ALTER SESSION statement and rerun the tablespace creation command.

You must set the db_create_file_dest parameter using the alter session or system command and re issue the statement.

```
Alter Session set db_create_file_dest ='d:\oracle\data\my_datafiles';
```

Question 49.

The TYPE clause used in the creation of an external table defines which access driver is used to convert the data when needed by SQL statements inside the database. Which type of access driver is supported for external tables?

- A. Import
- B. ORACLE LOADER
- C. Metadata API
- D. Direct path export.

Answer: B

Explanation:

When creating an external table, an additional TYPE option is specified that determines which access driver to use to access the data, with the default being ORACLE LOADER. The options stating that external tables are always read-only, external tables use an access type of ORACLE_LOADER by default, and creating an external table allocates no extents, but simply updates the data dictionary are correct.

Question 50.

To assist you in migrating your existing LONG columns to LOB data types, the ALTER TABLE syntax has been enhanced to support LONG to CLOB migration and LONG RAW to BLOB migration.

During this migration, if redo logging is enabled for the table or for the LOB being created, the ALTER TABLE MODIFY command temporarily _____ the space requirements.

- A. halves
- B. doubles
- C. triples
- D. quadruples

Answer: B

Explanation:

The ALTER TABLE MODIFY LONG->LOB statement copies the contents of the table into a new space, and frees the old space at the end of the operation. This temporarily doubles the space requirements. But the advantage is that after the transformation, the table will not have any embedded NULLs, so the performance of subsequent DMLs or queries is good.