PassQueen
a.k.a dumpsqueen.com

Oracle Database 12c Essentials

Oracle 1z0-497

Version Demo

Total Demo Questions: 18

Total Premium Questions: 189

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Database Architecture	36
Topic 2, Installing Oracle Database Software	5
Topic 3, Creating an Oracle Database	1
Topic 4, Managing Database Storage Structures	12
Topic 5, Administering User Security	7
Topic 6, Managing Schema Objects	9
Topic 7, Managing Data Concurrency	10
Topic 8, Managing Undo Data	4
Topic 9, Implementing Oracle Database Security	6
Topic 10, Performing Database Backups	9
Topic 11, Performing Database Recovery	12
Topic 12, Moving Data	9
Topic 13, Maintaining the Database	13
Topic 14, Managing Performance	7
Topic 15, Managing Resources	1
Topic 16, Configuring Oracle Net Services	15
Topic 17, Administering Multitenant Container Databases	31
Topic 18, Mix Questions	2
Total	189

QUESTION NO: 1

Examine the command:

sqlplus /nolog

In which scenario would this command execute successfully?

A. only from the server, if the database instance is not up

B. from either a client or server, even if the database instance is not up

C. from either a client or server, only if the database instance is up, and the command connects the session to the instance

D. from either a client or server, only if the database instance is up, but the command does not connect the session to the instance

ANSWER: B

Explanation:

`sqlplus /nolog` starts the SQL*Plus program without attempting to establish a database connection. SQL*Plus is a client utility, so it can be launched from a computer with an Oracle client installation as well as directly on a database server, provided the SQL*Plus executable and its required Oracle environment are available. The `/NOLOG` argument explicitly instructs SQL*Plus not to log on when it starts; it presents the `SQL>` prompt in a disconnected state instead.

Because no connection is made during startup, SQL*Plus does not need a running database instance, an available listener, a service name, or valid database credentials for this command itself to complete successfully. Those requirements become relevant only when a subsequent `CONNECT` command is issued. This behavior is useful when preparing a SQL*Plus session before deciding how to connect, running commands that do not require an active database connection, or connecting later using operating-system authentication or explicit credentials.

Oracle documents `/NOLOG` as the startup option that prevents SQL*Plus from connecting to a database: Oracle SQL*Plus User's Guide: Starting SQL*Plus. See also the Oracle SQL*Plus connection documentation.

QUESTION NO: 2

Examine the table:

Table name: PRODUCTS

Schema: SCOTT

Tablespace: USERS

Organization: Standard (Heap Organized)

This table contains 20 rows. In the data that has been entered, the PROD_ID column contains only single-digit values, PROD_NAME has a maximum of five characters, and PROD_RATE has four-digit values for all rows.

Which two statements are true regarding the modifications that can be performed to the structure of the PRODUCTS table?

A. PROD_NAME size can be increased from five to 10 characters.

B. PROD_ID can be changed from CHAR to NUMBER data type.

C. PROD_RATE size can be changed from NUMBER(4) to NUMBER(6, 2).

D. New columns can be added only if they have a NOT NULL constraint.

E. None of the existing columns can be dropped because they contain data.

ANSWER: A C

Explanation:

`PROD_NAME` size can be increased from five to 10 characters because Oracle permits a character column's declared maximum length to be increased when the table already contains rows. This is a metadata-compatible structural change: all existing values that fit within five characters also fit within the new ten-character limit. The heap-organized nature of the table does not impose a restriction on this change.

`PROD_RATE` size can be changed from `NUMBER(4)` to `NUMBER(6, 2)` because the new definition allows four digits to the left of the decimal point: precision six minus scale two leaves four integer digits. Therefore, the existing four-digit numeric values remain valid under the revised precision and scale. Oracle permits a numeric-column definition to be modified when the existing stored values conform to the new definition. These changes are issued with the `ALTER TABLE ... MODIFY` statement; Oracle validates the resulting column definition against existing data as necessary. Oracle documents the supported column-modification operations and their restrictions in the ALTER TABLE SQL Reference. The meaning of numeric precision and scale is documented in Oracle's SQL Data Types reference.

QUESTION NO: 3

Your customer is looking for a solution to address these downtime needs:

- Online addition and removal of disks should automatically rebalance data across new storage configurations while storage, database, and applications remain online.

- Dynamically add or remove clustered nodes without disrupting the database or applications.

- Automate the sizing and distribution of shared memory.

- Add a column with a default value without disrupting the database or applications.

- Perform rolling patch upgrades.

Which four would you recommend to address those needs?

A. Automatic Storage Management (ASM)

B. Online Data and Application change

C. Global Data Services (GDS)

D. Real Application Clusters (RAC)

E. Oracle Streams

F. Online Table Redefinition

G. Oracle Database Vault

ANSWER: A B D F

Explanation:

Automatic Storage Management (ASM) is appropriate because it provides volume-management capabilities for Oracle database files and automatically rebalances data when disks are added to or removed from an ASM disk group. This enables storage configuration changes while the database remains available. **Real Application Clusters (RAC)** supports a clustered database with multiple instances accessing one database, allowing capacity to be adjusted by adding or removing instances and supporting rolling maintenance procedures. RAC also works with Oracle's automatic memory-management capabilities to manage memory resources across participating instances.

Online Data and Application change addresses availability during application and schema evolution. Oracle supports online DDL capabilities, including efficient addition of columns with default values in applicable cases, so that many changes can be made without taking the application offline. For more substantial structural changes, **Online Table Redefinition** permits a table's definition to be changed while it remains available for queries and DML. It copies and synchronizes data during the redefinition process before the final metadata switch, minimizing the interruption required for the change.

These technologies complement one another: ASM handles online storage reconfiguration, RAC supplies clustered availability and rolling operations, and Oracle's online change features reduce downtime for database and application evolution. See Oracle's ASM overview and online table redefinition documentation.

QUESTION NO: 4

You are installing OracleDatabase 12c on a machine. When you run the installer; Oracle Universal installer (OUI) shows a message that says one of the product-specific prerequisite checks has failed:

Checking available swap space requirements...

Expected result: 1512MB

Actual Result: 1018MB

Check complete. The overall result of this check is: failed ««

Problem: The system does not have the required swap space.

What happens to the installation in this situation?

A. It can be continued.

B. It resizes the swap space automatically when you proceed further.

C. It can be continued, but the instance cannot be started without increasing swap space.

D. It shows a message saying one or more prerequisite checks have failed and the installation aborts.

ANSWER: A

Explanation:

It can be continued. Oracle Universal Installer performs prerequisite checks before installation to identify conditions that do not meet Oracle Database requirements, including available swap space. A failed prerequisite check is presented on the installer's prerequisite-check screen, where the administrator can generally choose to correct the condition and rerun the check or elect to continue despite the warning. OUI does not automatically alter operating-system swap configuration; that remains an administrative task. Thus, the failed check does not inherently terminate the installer session or make continuation impossible.

In a real deployment, the appropriate practice is to provision the required swap space before proceeding. The prerequisite value is derived from Oracle's operating-system installation requirements and is intended to ensure sufficient virtual-memory resources during installation and database operation. Continuing is an installer capability, not confirmation that the host configuration is supported or suitable for production use. Oracle's installation documentation describes reviewing prerequisite-check results and resolving failures where possible before continuing. See the Oracle Database 12c installation prerequisite documentation and the Oracle Universal Installer documentation.

QUESTION NO: 5

Identify the correct order of consolidation density.

A. multitenant< schema consolidation< dedicated database< VMs

B. schema consolidation< VMs< dedicated databases< multitenant

C. dedicated databases< VMs< schema consolidation< multitenant

D. VMs< dedicated databases< schema consolidation< multitenant

ANSWER: D

Explanation:

VMs< dedicated databases< schema consolidation< multitenant is the correct ascending order of consolidation density. Virtual-machine-based consolidation has the lowest density because each database workload is commonly isolated within its own guest operating system, retaining substantial operating-system and database-instance overhead. Dedicated databases increase density by placing separate database instances on shared physical infrastructure while preserving database-level isolation. Schema consolidation places multiple applications or tenants into schemas within one database, so they share an Oracle instance and its memory, background processes, and infrastructure. Oracle Multitenant provides the highest density in this progression because multiple pluggable databases reside in a single container database and share the container's instance-level resources while retaining strong administrative and operational separation. This architecture allows resources such as memory and background processes to be shared efficiently, and it supports rapid provisioning and centralized administration of many pluggable databases. Oracle describes a container database as containing zero, one, or many pluggable databases, with the multitenant architecture designed to simplify database consolidation. See the Oracle Database Concepts overview of the multitenant architecture and the Oracle Multitenant administration documentation.

QUESTION NO: 6

You are connected to the root container of a CDB. Which statement creates a common user?

A. CREATE USER APPUSR IDENTIFIED BY password CONTAINER=ALL;

B. CREATE USER C##APPUSR IDENTIFIED BY password CONTAINER=ALL;

C. CREATE COMMON USER APPUSR IDENTIFIED BY password;

D. CREATE USER C##APPUSR IDENTIFIED BY password CONTAINER=CURRENT;

ANSWER: B

Explanation:

`CREATE USER C##APPUSR IDENTIFIED BY password CONTAINER=ALL;` creates a common user in the root and makes that user known in every current and future container in the CDB. In an Oracle Database 12c CDB using the default common-user naming rules, a common user name must begin with the prefix `C##` or `c##`.

The `CONTAINER=ALL` clause specifies that the user is common. A local user is created in a PDB and is limited to that PDB. Oracle documents common and local user administration in the Oracle Multitenant Administration Guide.

QUESTION NO: 7

Which two statements are true about Oracle Data Pump Export and Import?

A. Data Pump jobs are performed by processes running on the database server.

B. A database directory object identifies the server location for dump files.

C. Dump files are always written to the client computer that starts the job.

D. Data Pump requires that the target database be in NOMOUNT state.

E. A Data Pump directory object is created by using an operating-system command.

ANSWER: A B

Explanation:

Data Pump Export and Import are server-based utilities. The database server processes perform the export and import work, including reading and writing dump files. A Data Pump job uses a database directory object that maps to an operating-system directory accessible to the Oracle Database server.

The directory object is specified by the `DIRECTORY` parameter or by a default directory object. Data Pump does not write dump files to the client computer merely because the command is started there. Appropriate privileges on the directory object are required. Oracle documents Data Pump architecture and directory-object use in the Oracle Data Pump Overview.

QUESTION NO: 8

Which two actions are performed by the Log Writer (LGWR) process?

A. It writes redo entries from the redo log buffer to online redo log files.

B. It writes redo for a committed transaction before reporting commit success.

C. It writes dirty buffers from the buffer cache to data files.

D. It archives filled online redo log files.

E. It removes inactive sessions after client-process failure.

ANSWER: A B

Explanation:

LGWR writes redo entries from the redo log buffer to the online redo log files. It writes redo at commit time so that a successful commit is protected in the redo logs before Oracle confirms completion to the user. LGWR also writes redo before DBWn writes dirty buffers to data files when required by the write-ahead logging rule.

LGWR does not write modified data blocks to data files; that is the responsibility of DBWn. It also does not archive completed redo log groups. In ARCHIVELOG mode, an ARCn process copies a filled online redo log file to an archived redo log destination after a log switch. See Oracle Database Concepts: Oracle Database Processes.

QUESTION NO: 9

Which two statements are true about RMAN backups?

A. A backup set can contain one or more backup pieces.

B. RMAN can create incremental backups.

C. Every RMAN backup is an operating-system copy of a data file.

D. RMAN cannot back up archived redo log files.

E. A recovery catalog is mandatory for all RMAN backups.

ANSWER: A B

Explanation:

An RMAN backup set is a proprietary RMAN backup format that can contain one or more backup pieces. RMAN can create backup sets for database files, archived redo logs, control files, and server parameter files. RMAN tracks backup metadata in the control file and, when configured, also in a recovery catalog.

RMAN can perform incremental backups. A level 0 incremental backup is the base for subsequent incremental backups, while higher-level incremental backups contain blocks changed since an appropriate earlier backup. RMAN backups are not necessarily operating-system copies of data files; image copies are the RMAN backup type that is a bit-for-bit copy of a single file. See the Oracle Database Backup and Recovery User's Guide: RMAN Backup Concepts.

QUESTION NO: 10

Examine the statements:

SQL> CREATE TABLESPACE MYTBS DATAFILE '/disk1/mytbs_l.dbf ' SIZE 10M;

SQL> CREATE TABLE mytab (id NUMBER, descr VARCHAR2 (100))

PARTITION BY RANGE(id) (

partition pi values loss than (100000) tablespacemytbs storage (initial lm),

partition p2 values less than (MAXVALUE) tablespacemytbs storage (initial 1m)

);

When inserting data, you get the error:

ORA-01688: unable to extend table SYS.MYTAB partition P2 by 128 in tablespace MYTBS

Which three operations would allow you to insert data without getting the error message?

A. Extend the logical volume by 10 MB where the data files are stored;
lvextend -L+10M /diskl

B. Add a data file:
SQL> ALTER TABLESPACE MYTBS ADD DATAFILE '/disk1/mytbs_2.dbf' SIZE 10M;

C. Move a partition to another tablespace:
SQL> CREATE TABLESPACE ADDITIONAL_TBS DATAFILE '/disk1/additional_tbs_1.dbf' SIZE 10M;
SQL> ALTER TABLE MYTAB MOVE PARTITION P2 TABLESPACE ADDITIONAL_TBS;

D. Resize the existing data file:
SQL> ALTER DATABASE DATAFILE '/disk1/mytbs_l.dbf' RESIZE 20M;

ANSWER: B C D

Explanation:

ORA-01688 indicates that Oracle cannot allocate the next required extent for partition `P2` in `MYTBS`. Adding a data file to `MYTBS` increases the tablespace's available storage and allows Oracle to allocate further extents for the partition. Resizing the existing data file has the same effect when the underlying file system or storage device has sufficient available capacity: the Oracle data file becomes larger, increasing the space managed by `MYTBS`.

Moving partition `P2` to `ADDITIONAL_TBS` is also a valid solution because the partition's segments and subsequent extent allocations are then stored in the target tablespace. The newly created target tablespace has a data file with available capacity, so inserts directed to the `MAXVALUE` range partition can continue. Oracle supports adding data files and resizing data files as tablespace capacity-management operations, and supports moving individual table partitions to another tablespace with `ALTER TABLE ... MOVE PARTITION`. See the Oracle Database Administrator's Guide: Managing Tablespaces and the Oracle SQL Language Reference: ALTER TABLE.

QUESTION NO: 11

A standby database is in mounted state and you are required to perform a failover to the physical standby. Which solution is best suited forthis scenario?

A. ALTER DATABASE RECOVER MANAGED STANDBY DATABASE statement with the FINISH keyword

B. ALTER DATABASE PREPARE TO SWITCHOVER and ALTER DATABASE COMMIT TO SWITCHOVER statements

C. ALTER DATABASE RECOVER TO STANDBY DATABASE statement with the FINISH keyword

D. D) ALTER DATABASE PREPARE SWITCHOVER and ALTER DATABASE COMMIT TO SWITCHOVER Statements

ANSWER: A

Explanation:

`ALTER DATABASE RECOVER MANAGED STANDBY DATABASE FINISH` is the appropriate statement to begin a physical-standby failover when the standby is mounted. The `FINISH` clause directs managed recovery to apply all redo that has already been received by the standby database and then complete recovery. This brings the physical standby as close as possible to the former primary's latest available state before it is converted to the primary role. Oracle documents this as the redo-application step in a failover procedure, including situations where the primary database is unavailable and a role transition cannot be performed as a switchover. After recovery has finished, the administrator activates the standby database with `ALTER DATABASE ACTIVATE STANDBY DATABASE`, then opens it for normal primary-database use. Thus, the statement with the `FINISH` keyword is best suited to the mounted physical standby at the recovery-completion stage of failover. This procedure helps minimize data loss to the redo that was successfully transported to and applied by the standby before activation. See Oracle's guidance on performing Data Guard role transitions and the ALTER DATABASE SQL reference.

QUESTION NO: 12

Which two statements are true about listeners?

A. Listeners use only the TCP/IP protocol.

B. Multiple listener processes can run simultaneously on a host.

C. Multiple database instances can be registered with a single listener,

D. Listener-related errors can be traced only at the administrative level.

E. Only one database instance can be registered with a single listener at any time.

ANSWER: B C

Explanation:

Multiple listener processes can run simultaneously on a host. Oracle Net permits a host to run more than one listener. Each listener is a separate process and is normally distinguished by its listener name and network endpoint configuration, such as a unique port number, protocol address, or both. This supports configurations where separate database environments, applications, or connection policies require independent listener endpoints.

Multiple database instances can be registered with a single listener, A single Oracle Net listener can accept connection requests for multiple database services and instances. Instances may register their services dynamically with the listener through the Local Listener or Remote Listener configuration, or they can be described by static service information in `listener.ora`. When a client requests a service, the listener uses its registered service information to direct the connection to the appropriate database instance. This architecture allows one listener endpoint to provide connectivity for several databases or instances on the same server or across configured network addresses.

Oracle documents listener configuration, multiple listener addresses, and service registration in the Oracle Net Services Administrator's Guide: Configuring and Administering the Listener and the Oracle Net Services concepts documentation.

QUESTION NO: 13

You create a CDB with four containers: the loot, the seed, and two pluggable databases PDB1 and PDB2. Which four items can reside In PDB1, but not in the root?

A. shared UNDO and default database temporary tablespace

B. Oracle-supplied metadata

C. shared Oracle-supplied data

D. CDB views providing information across PDBs

E. application tablespaces

F. local temporary tablespaces

G. local users and local roles

H. non-shared local metadata

ANSWER: E F G H

Explanation:

A PDB is a self-contained, portable collection of schemas, schema objects, and associated data files within a multitenant container database. Therefore, *application tablespaces* can be created in PDB1 to hold application data and are local to that PDB. PDB1 can also have *local temporary tablespaces*, enabling temporary segments for operations performed by users connected to that PDB without making those temporary tablespaces part of the root container.

Local users and local roles belong only to PDB1. Their identities, privileges, and role definitions are scoped to that individual PDB, allowing applications in separate PDBs to maintain independent user administration. A PDB also stores *non-shared local metadata*. This includes metadata for objects that are private to the PDB, such as locally created schemas and their objects. Oracle distinguishes this PDB-local metadata from the common metadata maintained for Oracle-supplied components in the root and made available across containers. Together, these features allow each PDB to operate as an isolated database environment while using the CDB architecture. See Oracle's Multitenant Architecture Overview and Managing Tablespaces in a CDB.

QUESTION NO: 14

Which command do you use to create a local role while you are connected as the local user?

A. create local role l_hr;

B. create role c##_hr container=current;

C. create role l_hr container=current;

D. create role l_hr container=all;

ANSWER: C

Explanation:

`create role l_hr container=current;` creates the role in the current container only. When the session is connected to a pluggable database, `CONTAINER=CURRENT` makes the role a local role of that PDB, rather than a common role available throughout the multitenant container database. A local user with the required `CREATE ROLE` system privilege can create and manage roles in that PDB. The role name `l_hr` is appropriate for a local role because local users and roles do not use the common-principal prefix required for common users and common roles. Explicitly specifying `CONTAINER=CURRENT` also makes the intended scope clear, even though this is the applicable local-container scope for the operation. Oracle documents the `CONTAINER` clause of `CREATE ROLE` as controlling whether a role is created in the current container or across all containers; local roles are created with the current-container setting. See the Oracle Database SQL Language Reference: CREATE ROLE and the Oracle Multitenant Administration Guide: Users and Roles in a CDB.

QUESTION NO: 15

When you start up a CDB, which operations are performed automatically and in what order?

a - The instance is started.

b - Control files are opened.

c - The root container is opened (redo logs and root data files).

d - Seed pluggable database is in READ ONLY mode.

e - Other PDBs are still in MOUNTED mode.

f - Triggers may fire if they exist to open other PDBs.

A. a, b, c, d, e, f

B. b, c, a, d, e, f

C. c, b, a, d, f, e

D. a, b, c, d

E. b, c, a, d

ANSWER: A

Explanation:

The correct startup order is **a, b, c, d, e, f**. Starting a container database first starts the Oracle instance, which allocates the required memory structures and starts the background processes. The database is then mounted by opening its control files. Once mounted, the root container is opened; this opens the root data files and online redo logs, making the root available for normal database operations.

As part of opening the container database, the seed pluggable database is automatically opened in read-only mode. The seed is a protected template used when creating new pluggable databases, so its read-only state is an intentional part of normal CDB startup. Other user-created pluggable databases do not automatically become open merely because the root opens; they remain mounted unless their saved state or configured startup behavior causes them to open. After the root is open, database startup triggers can execute. A trigger can therefore issue commands that open one or more PDBs, subject to the privileges and configuration of the trigger owner.

Oracle documents the CDB startup sequence and the automatic read-only opening of the seed PDB in its CDB startup and shutdown documentation. The SQL command behavior is also described in the STARTUP command reference.

QUESTION NO: 16

You are in the middle of transaction, updating a table. The machine on which the database is running reboots because of a power outage. This causes a database instance failure.

Which statement is true in this situation?

A. The online redo log files and archived redo log files are required to accomplish a recovery.

B. The uncommitted transaction is committed at the next startup of the database instance.

C. The uncommitted transaction is rolled back automatically at the next opening of the database.

D. The DBA has to perform a recovery on the database to recover the uncommitted transaction.

ANSWER: C

Explanation:

The uncommitted transaction is rolled back automatically at the next opening of the database. Following an instance failure, Oracle performs instance recovery automatically when the database is restarted. Instance recovery first applies the changes recorded in online redo to bring data blocks forward to a transactionally consistent point. Oracle then uses undo information to roll back all transactions that had not committed when the failure occurred. Consequently, the changes made by the interrupted transaction do not become permanent or visible as committed work after the database opens. This behavior preserves transaction atomicity: a transaction is either fully committed or its uncommitted effects are removed. Recovery work may begin during startup and remaining rollback work can, in some circumstances, continue through Oracle's transaction recovery mechanisms, but users do not need to manually recover an ordinary uncommitted transaction caused by an instance crash. Oracle documents that instance recovery is automatic and that uncommitted transactions are rolled back using undo data. See the Oracle Database Concepts description of instance recovery and the Oracle Database Administrator's Guide discussion of undo and transaction rollback.

QUESTION NO: 17

John is a CDB administrator, and Mike and Tim are PDB administrators of their respective databases- John takes a complete CDB backup by using RMAN>Backup Database. Mike and Tim may need point-in-time recovery of their PDBs. Which statement is true?

A. Recovery of individual PDBs is not possible, because John did not back up by using RMAN>Backup Pluggable database pdb1 pdb2. However, complete CDB recovery is still possible.

B. Mike and Tim need to individually take backups of their respective PDBs, because CDB-level backup does not include PDB data dictionary objects.

C. Mike and Tim do not need to worry, because recovery of individual PDBs is possible.

D. Recovery of individual PDBs is not possible at all.

ANSWER: C

Explanation:

Mike and Tim do not need to worry, because recovery of individual PDBs is possible. An RMAN `BACKUP DATABASE` operation performed while connected to the root of a container database backs up the database files required for the CDB, including the data files belonging to its PDBs. Therefore, the CDB-level backup provides the backup material from which an individual PDB can be restored and recovered; a separate PDB-only backup is not inherently required simply to make PDB point-in-time recovery possible.

Oracle RMAN supports PDB point-in-time recovery (PITR). The recovery procedure can target one or more PDBs with commands such as `RECOVER PLUGGABLE DATABASE ... UNTIL TIME`, restoring the required PDB data files from available backups and applying redo up to the requested point. Recovery is coordinated from the CDB root and has operational prerequisites, including an appropriate target time and the required archived redo logs. Oracle also manages the necessary CDB-level recovery context during the process, so a PDB can be recovered without recovering every PDB in the CDB to that earlier time. See Oracle's RMAN CDB and PDB backup and recovery documentation and its RMAN data repair concepts documentation.

QUESTION NO: 18

Which two statements are true about Oracle Database Resource Manager?

A. It can allocate CPU resources among consumer groups.

B. It can limit resources used by sessions assigned to a consumer group.

C. It replaces Oracle Net Services for client connection routing.

D. It stores archived redo log files in the fast recovery area.

E. It automatically creates a new database instance for each consumer group.

ANSWER: A B

Explanation:

Oracle Database Resource Manager can allocate CPU resources among consumer groups and can limit the resources used by sessions in those groups. A resource plan defines how resources are allocated, while consumer groups classify sessions that have similar processing requirements or priorities.

For example, a DBA can give interactive online users a greater CPU share than batch-processing sessions during business hours. Resource Manager operates within the database and does not replace operating-system resource controls or Oracle Real Application Clusters. Oracle documents resource plans, consumer groups, and CPU allocation in the Oracle Database Administrator's Guide.