

Oracle Database 19c: RAC, ASM, and Grid Infrastructure Administration

Oracle 1z0-078

Version Demo

Total Demo Questions: 9

Total Premium Questions: 93

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Installing and Configuring Oracle Grid Infrastructure	9
Topic 2, Managing Oracle Clusterware	24
Topic 3, Managing Oracle ASM	27
Topic 4, Managing Oracle RAC	33
Total	93

QUESTION NO: 1

Which two statements are true concerning the installation of Oracle Grid Infrastructure 19c Cluster? (Choose two.)

- A. The Grid Infrastructure Management Repository is only required when using ASM.
- B. Specifying the use of Grid Naming Service is optional.
- C. Normal Redundancy ASM Disk Groups used for Clusterware files will contain three voting files regardless of the number of nodes in the cluster.
- D. The Clusterware binaries may be installed without installing ASM binaries.
- E. High Redundancy ASM Disk Groups used for Clusterware files will contain two OCRs.

ANSWER: B C

Explanation:

Specifying the use of Grid Naming Service is optional. Grid Naming Service (GNS) provides dynamic name and address management for clusters, which is particularly useful where DHCP and dynamic VIP allocation are used. However, it is not a prerequisite for an Oracle Grid Infrastructure cluster installation; an administrator can instead configure the cluster using static network addresses and conventional DNS or hosts-file resolution. Oracle documents GNS as an optional Grid Infrastructure feature whose configuration depends on the site's networking design.

A Normal Redundancy ASM disk group selected for Oracle Clusterware files contains three voting files, regardless of the number of cluster nodes. Oracle Clusterware uses voting files to determine cluster membership and maintain quorum. When voting files are stored in ASM, their number is determined by the redundancy level of the ASM disk group, rather than by cluster size. Normal redundancy supplies the required three voting files and protects them through ASM's two-way mirroring architecture. This arrangement enables Clusterware to continue making quorum decisions when a voting-file failure occurs, subject to the available failure groups and disk-group health. See Oracle's [Grid Naming Service documentation](#) and [Oracle Clusterware management documentation](#).

QUESTION NO: 2

Which statement describes the usage of Dynamic Service Fallback in Oracle RAC 19c?

- A. Dynamic Service Fallback allows a service to relocate to an available instance when the service is stopped on a preferred instance.
- B. Dynamic Service Fallback allows a service to fall back to a preferred instance when the available instance performance falls below a specified threshold after it was failed over to an available instance.
- C. To enable Dynamic Service Fallback, you should specify the level of performance of a preferred instance for a service to be relocated to an available instance by setting the -fallback attribute of a service to "Low, Medium, or High".
- D. Dynamic Service Fallback allows a service to fall back to a preferred instance when it becomes available after it was failed over to an available instance.

ANSWER: D

Explanation:

Dynamic Service Fallback allows a service to fall back to a preferred instance when it becomes available after it was failed over to an available instance. In Oracle RAC, a service normally has preferred instances, where it is intended to run, and may have available instances, which can host it when the preferred placement cannot be used. When the final preferred instance hosting a service fails, Oracle Clusterware can start the service on an available instance to preserve client access. With Dynamic Service Fallback enabled, the service is automatically moved back to a preferred instance when that preferred instance is subsequently available again. This restores the intended service placement without requiring an administrator to manually relocate the service. The behavior is enabled through the service fallback setting, such as by using the SRVCTL service configuration option `-fallback YES`. Oracle documents this as a 19c service-management capability intended to

automate return to preferred service placement after failover. See the [Oracle RAC Administration and Deployment Guide](#) and the [Server Control \(SRVCTL\) Utility Reference](#).

QUESTION NO: 3

An ASM instance must start by using an ASM server parameter file stored in an ASM disk group.

No local text initialization parameter file is specified in the startup command.

How does Oracle Grid Infrastructure locate the ASM server parameter file?

- A. By reading the Grid Plug and Play (GPnP) profile
- B. By reading the Oracle Cluster Registry (OCR)
- C. By reading the database control file
- D. By reading the voting files

ANSWER: A

Explanation:

Oracle Grid Infrastructure uses the Grid Plug and Play (GPnP) profile to locate an ASM server parameter file stored in an ASM disk group. The GPnP profile contains the ASM discovery and configuration information needed during ASM startup.

The OCR is a cluster configuration repository, and the database control file contains database metadata. Neither is used as the bootstrap location for an ASM SPFILE.

QUESTION NO: 4

Which three statements are true regarding multitenant architecture for RAC databases? (Choose three.)

- A. One undo tablespace for each pdb is required.
- B. PDBs can have local temporary tablespaces.
- C. You can open one, several, or all PDBs on one, several, or all CDB instances.
- D. It requires the CDB to use shared undo management.
- E. One undo tablespace for each instance is required.
- F. PDBs must use shared temporary tablespaces.

ANSWER: B C E

Explanation:

In an Oracle RAC environment, each database instance requires its own undo tablespace. This is a fundamental RAC requirement because every instance performs transaction processing independently and must maintain its own undo data, even though all instances access the same database files. For a multitenant RAC database using local undo, the undo configuration is extended across containers as needed, but the per-instance undo requirement remains applicable. Oracle documents RAC undo administration in its [Oracle RAC Administration and Deployment Guide](#).

A PDB can use a local temporary tablespace. Local temporary tablespaces are associated with an individual PDB and enable temporary segments for that PDB to be managed separately from temporary space used by other containers. This capability is part of multitenant temporary-table-space administration, described in the [Oracle Multitenant Administrator's Guide](#).

Oracle RAC also allows PDB availability to be controlled on an instance-by-instance basis. A PDB may be opened on one instance, selected instances, or all CDB instances. This flexibility supports workload placement and services that are configured to run on particular RAC instances while preserving the shared-database architecture.

QUESTION NO: 5

A company expands a RAC cluster from three nodes to six nodes.

Clients use a Single Client Access Name (SCAN) to connect to database services.

Which statement is true about SCAN name resolution?

- A. A SCAN name resolves to up to three IP addresses regardless of cluster size.
- B. A SCAN name must resolve to one IP address for every cluster node.
- C. A SCAN name resolves only to the public IP address of the first cluster node.
- D. A SCAN name must resolve to the private interconnect addresses of all nodes.

ANSWER: A

Explanation:

A SCAN name resolves to up to three IP addresses, regardless of the number of nodes in the cluster. The SCAN listeners associated with those addresses provide a consistent client connection endpoint as nodes are added or removed.

A separate SCAN address is not required for every cluster node. Node VIP addresses remain node-specific and are different from SCAN addresses.

QUESTION NO: 6

An administrator adds two new ASM disks to disk group DATA and then removes an older disk from the same disk group.

Which two statements are true concerning the rebalance operations? (Choose two.)

- A. ASM automatically starts a rebalance when disks are added to or dropped from a disk group.
- B. A higher rebalance power value can increase the rate of rebalance work and resource consumption.
- C. The disk group must be dismounted on all ASM instances before any rebalance can begin.
- D. Adding disks changes the allocation unit size originally chosen for the disk group.
- E. A rebalance copies all database I/O through the ASM instance rather than directly to ASM disks.

ANSWER: A B

Explanation:

ASM automatically initiates rebalancing when disks are added to or dropped from a disk group. During the operation, ASM redistributes extents to reflect the changed disk-group configuration and to restore an appropriate distribution of mirrored extents where applicable.

The rebalance power setting controls the work rate assigned to the operation. A higher power value can complete the rebalance more quickly but can consume more system and storage resources. The disk group remains accessible during an online rebalance, subject to normal disk-group health and redundancy requirements.

QUESTION NO: 7

Which three statements are true about fully qualified file names in ASM? (Choose three.)

- A. They can be created only if OMF is enabled.
- B. They are system Generated at the time of ASM file creation.
- C. They can be used to reference existing ASM files.
- D. They are the only valid form that can be used for single file creation of a control file.
- E. An alias may be created for an existing file by referring to the fully qualified file name.

ANSWER: B C E

Explanation:

“They are system Generated at the time of ASM file creation,” “They can be used to reference existing ASM files,” and “An alias may be created for an existing file by referring to the fully qualified file name” are correct because a fully qualified ASM file name is the canonical, automatically generated name assigned when ASM creates a file. It includes the disk group, database name, file type, file number, and incarnation number, making it unique and suitable for identifying a specific existing ASM file. For example, the generated path identifies a particular data file or control file independently of any user-friendly alias.

Oracle documents that fully qualified names are generated by ASM and are used to refer to files that already exist. ASM aliases provide a more meaningful or convenient alternate name, but an alias maps to the underlying ASM file; it can be created by specifying that file's fully qualified name. This distinction is important operationally: the fully qualified name remains the file's immutable system identity, while aliases are optional namespace entries that can simplify administration and application configuration. See Oracle's [ASM Files, Directories, and Templates documentation](#) and the [ASM File Names documentation](#).

QUESTION NO: 8

Which statement is true concerning the installation of an Oracle Grid Infrastructure 19c patchset and its status during the installation?

- A. They can be applied in-place.
- B. When performing rolling patches, the vips for the node being patched are relocated to another node.
- C. Some Grid Infrastructure patchsets may not be installed in a rolling fashion.
- D. When performing rolling patches, crsctl query crs softwareversion always displays the earliest version of the software running anywhere in the cluster.

ANSWER: C

Explanation:

Some Grid Infrastructure patchsets may not be installed in a rolling fashion. Oracle Grid Infrastructure patching is classified as either rolling or nonrolling according to the specific patch and the components it modifies. A rolling installation is possible only when Oracle certifies that the patch can be applied node by node while the remaining cluster nodes continue to provide service. This minimizes cluster-wide downtime, but it is not a universal property of Grid Infrastructure maintenance. Patches that alter shared clusterware functionality, require synchronized changes across nodes, or are explicitly designated as nonrolling must be installed with the appropriate nonrolling procedure. In that case, the cluster resources and nodes are handled as required by the patch instructions so that every node ultimately runs the compatible patched Grid Infrastructure software. Administrators must therefore review the patch README and Oracle's patch documentation before planning maintenance, rather than assuming that every Grid Infrastructure patchset supports rolling application. Oracle documents the distinction between rolling and nonrolling Grid Infrastructure patching and directs administrators to use the method specified for the patch. See [Oracle Database 19c RAC Administration and Deployment Guide: Patching Oracle Grid Infrastructure](#) and [Oracle Grid Infrastructure Installation and Upgrade Guide: Automated Patching](#).

QUESTION NO: 9

A client connects to a RAC database by using a SCAN name and requests a service that is available on several instances.

Which statement correctly describes the role of the SCAN listener after the client connection reaches it?

- A.** It redirects the client to the local listener of an appropriate service instance.
- B.** It establishes every database session directly without using a local listener.
- C.** It connects the client only to an ASM instance on the SCAN listener node.
- D.** It requires clients to specify the public host name of every RAC node.

ANSWER: A

Explanation:

The SCAN listener accepts the initial client connection and, after determining an appropriate instance for the requested service, redirects the client to the local listener for that instance. The local listener then establishes the connection to the service handler.

A SCAN listener does not itself host database server processes for every connection, and it does not replace the local node listener. This redirection architecture allows clients to use a single SCAN name while Oracle RAC manages service placement across the cluster.