



# 1Z0-058<sup>Q&As</sup>

Oracle Real Application Clusters 11g Release 2 and Grid Infrastructure Administration

## Pass Oracle 1Z0-058 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4lead.com/1Z0-058.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





## QUESTION 1

You are managing a single-instance database and your company wants to convert the single instance database to an Oracle RAC database. You plan to use the rconfig utility to accomplish this task.

What are the prerequisites for converting the single-instance database to an Oracle RAC database?

- A. Oracle Clusterware and Oracle Database software are installed on all target nodes.
- B. The Oracle Database binary is enabled for Oracle RAC on all target nodes.
- C. The database being converted has been backed up successfully.
- D. Oracle Clusterware is stopped.
- E. Automatic Storage Management (ASM) must be enabled and the ASM disk should be available on all nodes.

Correct Answer: C

C.3.1 Prerequisites for Converting to Oracle RAC Databases Before you convert a single-instance database to an Oracle RAC database, ensure that the following conditions are met for each cluster node that you intend to make an Oracle RAC database node:

Oracle Clusterware 11g release 2 (11.2) is installed, configured, and running. Oracle RAC 11g release 2 (11.2) software is installed.

The Oracle binary is enabled for Oracle RAC.

Shared storage, either Oracle Cluster File System or Oracle ASM, is available and accessible from all nodes.

User equivalence exists for the oracle account.

If you intend to use Oracle Enterprise Manager, then all Oracle Enterprise Manager agents are configured and running, and are configured with cluster and host information.

You have backed up your existing database.

Oracle?Real Application Clusters Installation Guide 11g Release 2 (11.2) for Linux and UNIX

---

## QUESTION 2

Examine the Exhibit.

\*\*\* Exhibit is Missing \*\*\*

Which three statements are correct?

- A. Global Cache Block Access Latency indicates how many block transfers incurred a delay (busy) or an unexpected longer delay (congested).
- B. Global Cache Block Access Latency indicates how many block transfers received blocks cached by other instances as well as blocks read from disk.



- C. Global Cache Block Transfer Rate shows the total aggregated number of blocks received by all instances in the cluster over the interconnect.
- D. Global Cache Block Transfer Rate is essentially a profile of how much work is performed in the local buffer cache, rather than the portion of remote references and physical reads, which both have higher latencies.
- E. Global Cache Block Transfers and Physical Reads shows the percentage of read operations that retrieved blocks from disk, and the buffer cache of other instances using Cache Fusion.
- F. Global Cache Block Transfers and Physical Reads shows the aggregated number of data blocks received by all instances in the cluster over the interconnect.

Correct Answer: ACE

**Global Cache Block Access Latency** The global cache access latency represents the end-to-end elapsed time for a block request. The request is timed from when the request is initiated until it completes. Cache transfer indicates how many current and CR blocks per block class were received from remote instances, including how many transfers incurred a delay (busy) or an unexpected longer delay (congested).

If accessing a database block of any class does not locate a buffered copy in the local cache, a global cache operation is initiated. Before reading a block from disk, an attempt is made to find the block in the buffer cache of another instance. If the block is present in another instance, a version of the block may be shipped. Two different kinds of blocks are distinguished: current and consistent read blocks. The average block receive time represents the end-to-end elapsed time or latency for a block request.

**Global Cache Block Transfer Rate** The global cache block transfer rate shows the total aggregated number of data blocks received by all instances in the cluster by way of an interconnect. If a logical read fails to find a copy of the buffer in the local cache, it attempts to find the buffer in the database cache of a remote instance. If the block is present in another database, it is sent to the current instance.

**Global Cache Block Transfers and Physical Reads** The chart represents the percentage of logical reads that read data from the buffer cache of other instances via Direct Memory Access and from disk. It is essentially a profile of how much work is performed in the local buffer cache, rather than the portion of non-local references that incur some latency overhead.

**Active Sessions for Cluster Wait Class** The chart represents the active sessions for the cluster wait class. You can view the top modules and top SQL for a time period by selecting the time period in the active sessions chart. Oracle Enterprise Manager Online Help

---

### QUESTION 3

You notice that there is a very high percentage of wait time for the gc current split event in your RAC database that has frequent insert operations. Which two recommendation would you make to reduce this problem?

- A. shorter transactions
- B. using hash partitioned global indexes
- C. uniform and large extent sizes
- D. automatic segment space management
- E. smaller extent sizes
- F. increasing sequence cache sizes



Correct Answer: DF

---

#### QUESTION 4

Which four statements are true about ADVM interoperability?

- A. Using fdisk or similar disk utilities to partition ADVM-managed volumes is not supported
- B. On Linux platforms, the raw utility can be used to map ADVM volume block devices to raw volume devices.
- C. The creation of multipath devices over ADVM devices is not supported.
- D. You may create ASMLIB devices over ADVM devices to simplify volume management.
- E. ADVM does not support ASM storage contained in Exadata.
- F. ADVM volumes cannot be used as a boot device or a root file system.

Correct Answer: ACEF

Oracle Automatic Storage Management Cluster File System (Oracle ACFS) and Oracle ASM Dynamic Volume Manager (Oracle ADVM) extend Oracle ASM support to include database and application executables, database trace files,

database alert logs, application reports, BFILEs, and configuration files. Other supported files are video, audio, text, images, engineering drawings, and other general-purpose application file data.

Because of the fact that Oracle ADVM Volumes are technically spoken ASM files located on ASM Disk groups, and the fact that the Dynamic Volumes do not use the traditional device partitioning, it enables Oracle to extend some of the

ASM features to the ASM Clustered File Systems, which are created inside these ADVM Volumes, such as dynamic resizing or dynamically adding volumes. This makes ADVM and ACFS a far more flexible solution than traditional physical

devices.

Important Notes:

Partitioning of dynamic volumes (using fdisk or similar) is not supported Do not use raw to map ADVM volume block devices into raw volumes devices Do not create multipath devices over ADVM devices Do not create ASMLIB devices over

ADVM devices

Oracle ADVM supports all storage solutions supported for Oracle ASM with the exception of NFS and Exadata storage

ADVM volumes cannot be used as a boot device or a root file system

---

#### QUESTION 5

Which three fragments will complete this statement correctly? In a cluster environment, an ACFS volume\_\_\_\_\_.



- A. Will be automatically mounted by a node on reboot by default
- B. Must be manually mounted after a node reboot
- C. Will be automatically mounted by a node if it is defined as cluster stack startup if it is included in the ACFS mount registry.
- D. Will be automatically mounted to all node if it is defined as cluster resource when dependent cluster resources requires access
- E. Will be automatically mounted to all node in the cluster when the file system is registered
- F. Must be mounted before it can be registered

Correct Answer: ACE

The Oracle ACFS mount registry supports both Oracle Restart and Oracle Grid Infra- structure cluster configurations. File systems that are to be mounted persistently (across reboots) can be registered with the Oracle ACFS mount registry. In cluster configura- tions, registered Oracle ACFS file systems are automatically mounted by the mount registry, similar to a clusterwide mount table. By default, an Oracle ACFS file system that is inserted into the cluster mount registry is automatically mounted on all cluster members, including cluster members that are added after the registry addition. The Oracle ACFS registry resource actions are designed to automatically mount a file system only one time for each Oracle Grid Infrastructure initialization to avoid potential conflicts with administrative actions to dismount a given file system. Oracle Automatic Storage Management Administrator\\'s Guide

---

## QUESTION 6

Which two statements are true regarding the Active Session History (ASH) reports for RAC?

- A. They provide details about Oracle databases for all current sessions, and history of past session all RAC nodes.
- B. They provide statistics about Oracle databases for the active sessions on all the RAC nodes.
- C. They report on data captured for active sessions. The volume of data is directly related to the work being performed by sessions.
- D. They report on data captured for active sessions. The volume of data is directly related to the number of sessions on the system.

Correct Answer: AC

ASH report statistics provide details about Oracle Database session activity. Oracle Database re- cords information about active sessions for all active Oracle RAC instances and stores this data in the System Global Area (SGA). Any session that is connected to the database and using CPU is considered an active session. The exception to this is sessions that are waiting for an event that belongs to the idle wait class. ASH reports present a manageable set of data by capturing only in- formation about active sessions. The amount of the data is directly related to the work being performed, rather than the number of sessions allowed on the system. Oracle Real Application Clusters Administration and Deployment Guide

---

## QUESTION 7

Examine the set of commands given below:



```
$ echo $ORACLE_SID
RACDB_1
$ sqlplus / as sysdba
```

```
SQL> SELECT * FROM V$ACTIVE_INSTANCES;
INST_NUMBER INST_NAME
```

```
-----
1          gr7597:RACDB_1
2          gr7602:RACDB_2
3          gr7633:RACDB_3
```

```
SQL>ALTER SYSTEM SET TRACE_ENABLED=TRUE SCOPE=MEMORY sid='*'
SQL>ALTER SYSTEM SET TRACE_ENABLED=FALSE SCOPE=MEMORY sid='RACDB_1';
```

Which statement is true in this scenario?

- A. All the instances of the RACDB will have the TRACE\_ENABLED parameter with the value TRUE because this parameter must have the same value on all the instances.
- B. All the instances of the RACDB will have the TRACE\_ENABLED parameter with the value TRUE except the RACDB\_1 instance, which will have the value FALSE for this parameter.
- C. All the instances of the RACDB will have the TRACE\_ENABLED parameter with the value TRUE because the first ALTER statement has higher precedence.
- D. All the instances of the RACDB will have the TRACE\_ENABLED parameter with the value FALSE because the second ALTER statement has higher precedence.

Correct Answer: B

Section: (none)

SPFILE Parameter Values and RAC

- You can change parameter settings using the ALTER SYSTEM SET command from any instance:  
ALTER SYSTEM SET <pname> SCOPE=MEMORY sid='<sid|\*>';
- SPFILE entries such as:
  - \*.<pname> apply to all instances
  - <sid>.<pname> apply only to <sid>
  - <sid>.<pname> takes precedence over \*.<pname>
- Use current or future \*.<pname> settings for <sid>:  
ALTER SYSTEM RESET <pname> SCOPE=MEMORY sid='<sid>';
- Remove an entry from your SPFILE:  
ALTER SYSTEM RESET <pname> SCOPE=SPFILE sid='<sid|\*>';

## QUESTION 8





Which command will set full debugging for the communications and resource management components of the CRS module on nodes host01 and host03?

- A. `srvctl set log -m CRS -c CRSCTI 5, CRSCOMM 5 -m host01,host03`
- B. `ocrconfig set log CRS "CRSRTI, CRSCOMM: 5" -nodelist host01,host03`
- C. `crsctl set log CRS "CRSRTI: 5, CRSCOMM: 5" -nodelist host01,host03`
- D. `crsctl lsmodules crs set log CRSRTX : 5, CRSCOMM: 5 -nodelist host01, host03`

Correct Answer: C

### Dynamic Debugging

This section includes the following CRSCTL commands that aid in debugging:

`crsctl set log`

Use the `crsctl set log` command to set log levels for Oracle Clusterware.

### Syntax

`crsctl set log {[crs | css | evm "component_name=log_level, [...]"] | [all=log_level]}` You can also set log levels for the agents of specific resources, as follows:

`crsctl set log res "resource_name=log_level, [...]"`

**Table E-44 Oracle Clusterware (CRS) Modules and Functions**

Module	Descriptions
CRSUI	User interface module
CRSCOMM	Communication module
CRSRTI	Resource management module
CRSMAIN	Main module driver
CRSPLACE	CRS placement module
CRSAPP	CRS application
CRSRES	CRS resources
CRSOCR	Oracle Cluster Registry interface
CRSTIMER	Various timers related to CRS
CRSEVT	CRS EVM/event interface module
CRSD	CRS daemon

Oracle Clusterware Administration and Deployment Guide 11g Release 2 (11.2)

### QUESTION 9

The original cluster on RACNODE1, RACNODE 2, RACNODE 3, and RACNODE4 had the Oracle Grid Infrastructure installed to support Grid Naming Service (GNS). What should be done to verify that the two new nodes called RACNODE5, RACNODE6 are physically connected?

- A. `cluvfy stage -post crsinst -n RACNODE5,RACNODE6 -verbose`



- B. cluvfy stage -post nodeadd -n RACNODE5,RACNODE6 -verbose
- C. cluvfy stage -post hacfg -verbose
- D. cluvfy stage -post nodeadd -n all verbose

Correct Answer: B

Completing OUI Silent Node Addition Perform integrity checks on the cluster. [grid@host01]\$ cluvfy stage post nodeadd n host03 -verbose D60488GC11 Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 4 - 8

## QUESTION 10

Your four-node cluster was originally purchased, installed, and configured three years ago. You recently added another four nodes to the cluster.

Now you want to remove two of the older nodes that are still accessible to be redeployed elsewhere in the data center. Which two are true regarding the procedure for removing one or more cluster nodes?

- A. The procedure requires that all commands be invoked from one of the surviving cluster nodes.
- B. All commands are run as root regardless of which nodes are used to invoke them.
- C. The procedure requires that some commands be invoked on the node or nodes to be removed and that some be invoked from all surviving cluster nodes.
- D. The procedure requires that some commands be invoked on the node or nodes to be removed and that some be invoked from one surviving cluster node.
- E. Some commands require that the name of the node or nodes to be removed are passed as arguments, and some commands require the name of existing nodes to be passed.

Correct Answer: DE

If you are deleting multiple nodes, then run the rootcrs.pl script on each node that you are deleting. From any node that you are not deleting, run the following command from the Grid\_home/bin directory as root to delete the node from the cluster:

```
# crsctl delete node -n node_to_be_deleted
```

On the node you want to delete, run the following command as the user that installed Oracle Clusterware from the Grid\_home/oui/bin directory where node\_to\_be\_deleted is the name of the node that you are deleting:

```
$ ./runInstaller -updateNodeList ORACLE_HOME=Grid_home "CLUSTER_NODES= {node_to_be_deleted}"  
CRS=TRUE -silent -local
```

On the node that you are deleting, depending on whether you have a shared or local Oracle home, complete one of the following procedures as the user that installed Oracle Clusterware:

For a local home, deinstall the Oracle Clusterware home from the node that you want to delete, as follows, by running the following command, where Grid\_home is the path defined for the Oracle Clusterware home:

On any node other than the node you are deleting, run the following command from the Grid\_home /oui/bin directory where remaining\_nodes\_list is a comma-delimited list of the nodes that are going to remain part of your cluster:





```
$. /runInstaller -updateNodeList ORACLE_HOME=Grid_home "CLUSTER_NODES= {remaining_nodes_list}"  
CRS=TRUE -silent
```

Run the following CVU command to verify that the specified nodes have been successfully deleted from the cluster:

```
$. cluvfy stage -post nodedel -n node_list [-verbose] Oracle?Clusterware Administration and Deployment Guide
```

---

### QUESTION 11

When creating an Oracle Cluster database using DBCA the "Memory size (SGA and PGA)" field is supplied on value of 2000 MB. Identify the default block Size used for the database.

- A. 2 KB
- B. 4 KB
- C. 8 KB
- D. 16 KB
- E. 32 KB

Correct Answer: C

Type of Limit Limit Value Minimum 2k. Must be a multiple of operating system physical block size Maximum Operating system dependent, but never more than 32 KB DBCA tab sizing In this tab, you specify the smallest block size and the maximum number of operating system user processes that can simultaneously connect to the database. In the Block Size list, enter the size in bytes or accept the default. Oracle Database data is stored in these blocks. One data block corresponds to a specific number of bytes of physical space on disk. While using pre-defined templates, this field is not enabled since the database will be created with the default block size of 8 KB. But while using the custom option, you can change block size. Selecting a block size other than the default 8 KB value requires advanced knowledge and should only be done when absolutely required. Oracle?Database 2 Day DBA

---

### QUESTION 12

A policy-managed RAC database that hosts the ACCOUNTS service has only three instances running rather than the usual four on a six-node cluster. Five of the six cluster nodes are active and the sixth (RACNODE6) is down for maintenance. You decide to check the state of the servers in the server pools to see if there are problems with the OLTP pool to which the ACCOUNTS service has been assigned. You used the `ctscctl status server -f` command and see that the RACNODE3 node has STATE=VISIBLE.

What is true about this situation?

- A. The Cluster Synchronization Services Daemon (CSSD) is running RACNODE3 is considered to be part of the cluster, but the Cluster Ready Services Daemon (CRSD) is currently not running.
- B. Both the Cluster Ready Services Daemon (CRSD) and the Cluster Synchronization Services Daemon (CSSD) are running on RACNODE3, but RACNODE3 is currently being moved to another server pool.
- C. The Cluster Ready Services Daemon (CRSD) is running and RACNODE3 is considered to be part of the cluster, but the Cluster Synchronization Services Daemon (CSSD) is currently not running.
- D. Neither the Cluster Ready Services Daemon (CRSD) nor the Cluster Synchronization Services Daemon (CSSD) is



running on RACNODE3.

Correct Answer: A

STATE	A server can be in one of the following states:
ONLINE	The server is a member of the cluster and is available for resource placement.
OFFLINE	The server is not currently a member of the cluster. Subsequently, it is not available for resource placement.
JOINING	When a server joins a cluster, Oracle Clusterware processes the server to ensure that it is valid for resource placement. Oracle Clusterware also checks the state of resources configured to run on the server. Once the validity of the server and the state of the resources are determined, the server transitions out of this state.
LEAVING	When a planned shutdown for a server begins, the state of the server transitions to LEAVING, making it unavailable for resource placement.
VISIBLE	Servers that have Oracle Clusterware running, but not the Cluster Ready Services daemon (crsd), are put into the VISIBLE state. This usually indicates an intermittent issue or failure and Oracle Clusterware trying to recover (restart) the daemon. Oracle Clusterware cannot manage resources on servers while the servers are in this state.
RECONFIGURING	When servers move between server pools due to server pool reconfiguration, a server is placed into this state if resources that ran on it in the current server pool must be stopped and relocated. This happens because resources running on the server may not be configured to run in the server pool to which the server is moving. As soon as the resources are successfully relocated, the server is put back into the ONLINE state.

Use the `crsctl status server` command to obtain server

Oracle?Clusterware Administration and Deployment Guide 11g Release 2 (11.2)

### QUESTION 13

Which two statements are true about ACFS snapshots?

- A. They can be created for ACFS file systems only if the ASM disk group hosting the ADVM volume file used by the file system has free space available.
- B. They can be created for ACFS file systems only if the ADVM volume file used by the file system has free space available.
- C. They can be created only if the ASM disk group hosting the ADVM volume used by the file system has no other ASM files contained in the disk group.



- D. They can be created when ACFS is used both on clusters and on stand-alone servers.
- E. They are accessible only on the cluster node that was used when creating the snapshot.

Correct Answer: BD

About Oracle ACFS Snapshots Oracle ACFS snapshot storage is maintained within the file system, eliminating the management of separate storage pools for file systems and snapshots. Oracle ACFS file systems can be dynamically resized to accommodate additional file and snapshot storage requirements. Oracle Automatic Storage Management Administrator's Guide 11g Release 2 (11.2)

#### QUESTION 14

How does the Oracle Grid Infrastructure administrator determine the location of the Oracle Clusterware voting disk?

- A. Run `cat /etc/oracle/vote.loc` from any node.
- B. Run `srvctl query css votedisk` from any node.
- C. Run `crsctl query css votedisk` from any node.
- D. Run `select name, path from v$votedisk` from any RAC database instance of any database on the cluster.

Correct Answer: C

Determining the Location of Oracle Clusterware Configuration Files You can determine the location of the voting disk by using the `crsctl query css votedisk` command on any node. D60488GC11 Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 3 7

#### QUESTION 15

Which two statements are true regarding the Automatic Workload Repository (AWR) in a RAC environment?

- A. The AWR includes time model statistics based on time usage for activities displayed in the `V$SYS_TIME_MODEL` and `V$SESS_TIME_MODEL` views.
- B. The AWR is controlled by the `statistics_level` initialization parameter and it must be set to `TYPICAL` or `BASIC`.
- C. The Manageability Monitor Processes (MMON) process gathers statistics every hour from every instance and creates an AWR snapshot and stores it in the SGA.
- D. The MMON process gathers statistics from its own instance and kicks off statistics collection from other instances every hour from other instances and creates an AWR snapshot which is written to the `sysaux` tablespace.

Correct Answer: AD

The Automatic Workload Repository (AWR) collects, processes, and maintains performance statistics for problem detection and self-tuning purposes. This data is both in memory and stored in the database. The gathered data can be displayed in both reports and views. The statistics collected and processed by AWR include: Object statistics that determine both access and usage statistics of database segments Time model statistics based on time usage for activities, displayed in the `V$SYS_TIME_MODEL` and `V$SESS_TIME_MODEL` views Some of the system and session statistics collected in the `V$SYSSTAT` and `V$SESSTAT` views SQL statements that are producing the highest load on the system, based on criteria such as elapsed time and CPU time Active Session History (ASH) statistics, representing



the history of recent sessions Activity MMON (Memory Monitor) is a background process that gathers memory statistics (snapshots) stores this information in the AWR (automatic workload repository). MMON is also responsible for issuing alerts for metrics that exceed their thresholds. Oracle Database Performance Tuning Guide

[1Z0-058 PDF Dumps](#)

[1Z0-058 VCE Dumps](#)

[1Z0-058 Practice Test](#)



To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

## Try our product !

100% Guaranteed Success

100% Money Back Guarantee

365 Days Free Update

Instant Download After Purchase

24x7 Customer Support

Average 99.9% Success Rate

More than 800,000 Satisfied Customers Worldwide

Multi-Platform capabilities - [Windows](#), [Mac](#), [Android](#), [iPhone](#), [iPod](#), [iPad](#), [Kindle](#)

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications. You can view Vendor list of All Certification Exams offered:

<https://www.pass4lead.com/allproducts>

## Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 <p><b>One Year Free Update</b> Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 <p><b>Money Back Guarantee</b> To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 <p><b>Security &amp; Privacy</b> We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information &amp; peace of mind.</p>
---	---	--

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © pass4lead, All Rights Reserved.