

Oracle

Exam 1z0-058

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

Version: 8.8

[Total Questions: 139]

Question No : 1

Various clients can access and manipulate ASM files. Which two statements are true?

- A.** The DBMS_FILE_TRANSFER.COPY_FILE procedure can move a database file from one ASM to another ASM, but not to an operating system file system.
- B.** The ASMCMD cp command can move database files from a file system to ASM, but not from ASM to ASM.
- C.** The SQL*Plus command ALTER DISKGROUP orcl MOVE '+DATA/orcl/example01.dbf' to '+OLDDATA/orcl/example01.dbf' can move the example01 data file to a different diskgroup.
- D.** The DBMS_FILE_TRANSFER.GET_FILE procedure reads an ASM file from a remote machine and makes a local copy on an ASM or a file system.
- E.** The ASMCMD rm command will delete ASM files and directories, but not database files on an operating system file system.

Answer: D,E

Explanation:

DBMS_FILE_TRANSFER

COPY_FILE Procedure

This procedure reads a file from a source directory and creates a copy of it in a destination directory. The source and destination directories can both be in a local file system, or both be in an Automatic Storage Management (ASM) disk group, or between local file system and ASM with copying in either direction.

You can copy any type of file to and from a local file system. However, you can copy only database files (such as datafiles, tempfiles, controlfiles, and so on) to and from an ASM disk group.

GET_FILE Procedure

This procedure contacts a remote database to read a remote file and then creates a copy of the file in the local file system or ASM. The file that is copied is the source file, and the new file that results from the copy is the destination file. The destination file is not closed until the procedure completes successfully.

Examples

```
CREATE OR REPLACE DIRECTORY df AS '+datafile' ;  
GRANT WRITE ON DIRECTORY df TO "user";  
CREATE DIRECTORY DSK_FILES AS "^t_work^";  
GRANT WRITE ON DIRECTORY dsk_files TO "user";
```

-- assumes that dbs2 link has been created and we are connected to the instance.

-- dbs2 could be a loopback or point to another instance.

BEGIN

-- asm file to an os file

-- get an asm file from dbs1.asm/a1 to dbs2.^t_work^/oa5.dat

```
DBMS_FILE_TRANSFER.GET_FILE ( 'df' , 'a1' , 'dbs1', 'dsk_files' , 'oa5.dat' );
```

-- os file to an os file

-- get an os file from dbs1.^t_work^/a2.dat to dbs2.^t_work^/a2back.dat

```
DBMS_FILE_TRANSFER.GET_FILE ( 'dsk_files' , 'a2.dat' , 'dbs1', 'dsk_files' , 'a2back.dat' );
```

END ;

/

Oracle® Database PL/SQL Packages and Types Reference

11g Release 2 (11.2)

ASMCMD

cp

Purpose

Enables you to copy files between Oracle ASM disk groups and between a disk group and the operating system.

You can use the cp command to:

Copy files from a disk group to the operating system

Copy files from a disk group to a disk group

Copy files from the operating system to a disk group

rm

Purpose

Deletes the specified Oracle ASM files and directories.

Oracle® Automatic Storage Management Administrator's Guide

11g Release 2 (11.2)

Question No : 2

: 10

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

ADV volume file used by the file system has free space available.

- B.** They can be created for ACFS file systems only if the ADV volume file used by the file system has free space available.
- C.** They can be created only if the ASM disk group hosting the ADV 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.

Answer: B,D

Explanation:

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 No : 3



You are managing a policy-managed three-instance RAC database. You ran database ADDM for the database and noticed gc current block congested and gc cr block congested waits. What are two possible reasons for these wait events?

- A.** The wait events indicate a delay in processing has occurred in the Global Cache Services (GCS), which is usually caused by high load.
- B.** The wait times indicate that the blocks must wait after initiating a gc block request, for the round trip from the start of the wait until the blocks arrive.
- C.** The wait events indicate that there is block contention resulting in multiple requests for access to local blocks.
- D.** The wait events indicate that the local instance making the request for current or consistent read blocks was waiting for logical I/O from its own buffer cache at the same time.

Answer: A,B

Explanation: Load-Related Wait Events

The main wait events for load-related waits are:

-  gc current block congested
-  gc cr block congested

The load-related wait events indicate that a delay in processing has occurred in the GCS,

which is usually caused by high load, CPU saturation and would have to be solved by additional CPUs, load-balancing, off loading processing to different times or a new cluster node. For the events mentioned, the wait time encompasses the entire round trip from the time a session starts to wait after initiating a block request until the block arrives
Oracle Real Application Clusters Administration and Deployment Guide

Question No : 4

You have an ASM cluster and you want to terminate all the ASM instances. Assume that the database Instances supported by the ASM are already shut down. What must you know to select the appropriate command for this operation?

- A. Are the OCR and voting disk files stored In ASM diskgroups?
- B. Are the SYSTEM and SYSAUS tablespaces stored In ASM diskgroups?
- C. Is the undo tablespace stored in ASM diskgroups?
- D. Are the redo log files stored in ASM diskgroups?

Answer: A

Explanation: Oracle strongly recommends that you shut down all database instances that use the Oracle ASM instance and dismount all file systems mounted on Oracle ASM Dynamic Volume Manager (Oracle ADVM) volumes before attempting to shut down the Oracle ASM instance.

If Oracle Cluster Registry (OCR) or voting files are stored in a disk group, the disk group can only be dismounted by shutting down the Oracle ASM instance as part of shutting down the clusterware on a node. To shut down the clusterware, run `crsctl stop crs`.

Oracle Automatic Storage Management Administrator's Guide

Question No : 5

You have two administrator-defined server pools on your eight-node cluster called OLTP and DSS.

Hosts RACNODE3, RACNODE4, and RACNODE5 are currently assigned to the DSS Pool.

Hosts RACNODE6, RACNODE7, and RACNODE8 are assigned to the OLTP Pool.

Hosts RACNODE1 and RACNODE2 are assigned to the Generic pool.

You are patching the Oracle Grid Infrastructure in a rolling fashion for your cluster and you have completed patching nodes RACNODE3, RACNODE4, RACNODE5, and RACNODE6, but you have not patched nodes RACNODE1 and RACNODE2.

While examining the status of RACNODE2 software, you get this output:

```
$ crsctl query crs softwareversion
```

```
Oracle Cluster-ware version on node [RACNODE2] is [11.2.0.2.0]
```

```
$ crsctl query crs activeversion
```

```
Oracle Clusterware active version on node [RACNODE2] is [11.2.0.1.0]
```

Which two statements describe the reasons for the active versions on the nodes of the cluster?

- A.** The active version is 11.2.0.2.0 on RACNODE3, RACNODE4, and RACNODE5 because all the nodes in the DSS server pool have the same installed version.
- B.** The active version is 11.2.0.1.0 on RACNODE6, RACNODE7, and RACNODE8 because some nodes in the cluster still have version 11.2.0.1.0 installed.
- C.** The active version is 11.2.0.1.0 on RACNODE6, RACNODE7, and RACNODE8 because some nodes in the OLTP Pool still have version 11.2.0.1.0 installed.
- D.** The active version is 11.2.0.1.0 on RACNODE3, RACNODE4, and RACNODE5 because some nodes in the cluster still have version 11.2.0.1.0 installed.

Answer: B,D

Explanation:

```
crsctl query crs softwareversion
```

Use the `crsctl query crs softwareversion` command to display latest version of the software that has been successfully started on the specified node.

```
crsctl query crs activeversion
```

Use the `crsctl query crs activeversion` command to display the active version of the Oracle Clusterware software running in the cluster. During a rolling upgrade, however, the active version is not advanced until the upgrade is finished across the cluster, until which time the cluster operates at the pre-upgrade version.

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

Question No : 6

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

Answer: B

Explanation:

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 No : 7

Which command will resize an ACFS file system on the DATA diskgroup?

- The ADVM volume is named VOL1.
- The mount point is >/u01/app/oracle/afsmounts/vol1.
- The new, larger size is 20G.

- A. ASMCMD [+] > volsize -G data -s 20G VOL1
- B. /sbin/acfsutil size 20G /u01/app/oracle/afsmounts/vol1
- C. SQL> ALTER DISKGROUP data RESIZE VOLUME vol1 SIZE 20G;
- D. None, an ACFS file system cannot be resized.

E. fsadm /u01/app/oracle/afsmounts/vol1 20G

Answer: B

Explanation:

acfsutil size

Purpose

Resizes an Oracle ACFS file system.

Syntax and Description

acfsutil size -h

acfsutil size [+|-]n[K|M|G|T|P] [device] mount_point

acfsutil size -h displays help text and exits.

Oracle® Automatic Storage Management Administrator's Guide

11g Release 2 (11.2)

ALTER DISKGROUP

RESIZE VOLUME Clause

Use this clause to change the size of an existing Oracle ADVM volume. In an Oracle ASM cluster, the new size is propagated to all nodes. If an Oracle Automatic Storage Management File System (ACFS) exists on the volume, then you must use the acfsutil size command instead of the ALTER DISKGROUP statement.

Oracle® Database SQL Language Reference

11g Release 2 (11.2)

Question No : 8

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

Answer: C

Explanation: Type of LimitLimit Value

Minimum2k.Must be a multiple of operating system physical block size

MaximumOperating 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 No : 9

You have configured your eight-node cluster to use GNS. The network administrator has established delegated subdomain for the Cluster which is MYCLUSTER.EXAMPLE.COM.

DHCP has been configured so that the cluster now manages IP addresses within the cluster.

Select three responses that describe the VIPs that will exist in this configuration.

- A. 3 GNS VIPs
- B. 8 Node VIPs
- C. 3 SCAN VIPs
- D. 1 GNS VIP
- E. 3 Node VIPs

Answer: B,C,D

Explanation: Implementing GNS

If you use GNS, then you must specify a static IP address for the GNS VIP address, and delegate a subdomain to be delegated to that static GNS VIP address.

Dynamic IP address assignment using Oracle Grid Naming Service (GNS)

If you select this option, then network administrators assign static IP address for the physical host name and dynamically allocated IPs for the Oracle Clusterware managed VIP addresses. In this case, IP addresses for the VIPs are assigned by a DHCP and resolved using a multicast domain name server configured as part of Oracle Clusterware within the cluster. If you plan to use GNS, then you must have the following:

- ✍ A DHCP service running on the public network for the cluster
- ✍ Enough addresses on the DHCP to provide 1 IP address for each node's virtual IP,

and 3 IP addresses for the cluster used by the Single Client Access Name (SCAN) for the cluster

Oracle Grid Infrastructure Installation Guide

Question No : 10

Which three statements are true about the logging or tracing of Java-based Oracle tools?

- A. By default, Oracle logs actions for DBCA and DBUA.
- B. For tools such as SRVCTL and CLUVFY, tracing can be enabled by setting `srvm trace = true` your shell environment before executing the commands.
- C. The `srvm_trace` environment variable can be set only in the user's shell environment file.
- D. To disable tracing for Java-based tools, unset the `srvm_trace` variable: `export srvm_trace=`

Answer: A,B,D

Explanation: Enabling Tracing for Java-Based Tools and Utilities in Oracle RAC

All Java-based tools and utilities that are available in Oracle RAC are called by executing scripts of the same name as the tool or utility. This includes the Cluster Verification Utility (CVU), Database Configuration Assistant (DBCA), the Net Configuration Assistant (NETCA), Server Control Utility (SRVCTL), and the Global Services Daemon (GSD). For example to run DBCA, enter the command `dbca`. By default, Oracle Database enables traces for DBCA and the Database Upgrade Assistant (DBUA). For the CVU, GSDCTL, and SRVCTL, you can set the `SRVM_TRACE` environment variable to `TRUE` to make Oracle Database generate traces. Oracle Database writes traces to log files. For example, Oracle Database writes traces to log files in `Oracle home/cfgtoollogs/dbca` and `Oracle home/cfgtoollogs/dbua` for DBCA and DBUA, respectively.

Oracle® Real Application Clusters Administration and Deployment Guide

Question No : 11

In an ASM instance, there are several parameters that have meanings specific to the ASM. Which parameter is mandatory?