



Free Questions for 1Z0-902 by dumpshq

Shared by Bowman on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Which two Exadata security features would you configure to control which databases can access which griddisks, when multiple databases share space on your storage servers in an unpartitioned storage grid?

Options:

- A- Exadata storage realms using ASM-scoped security mode
- B- File permissions on the griddisks in each database server
- C- Using EXADCLI instead of CELLCLI to create the griddisks
- D- File permissions on the griddisks in each storage server
- E- Exadata storage realms using database-scoped security mode

Answer:

A, E

Explanation:

Exadata storage realms using ASM-scoped security mode allows you to control which database can access which griddisks by creating realms at the ASM level, thereby allowing you to limit the access of each database to the griddisks that it needs to access. This is covered in section 3.15.2 of the Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book and Resources [1].

Exadata storage realms using database-scoped security mode allows you to control which database can access which griddisks by creating realms at the database level[1], thereby allowing you to limit the access of each database to the griddisks that it needs to access. This is covered in section 3.15.3 of the Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book and Resources [2].

[1]<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/x9m/exad-implementing-database-machine-x9m.pdf>[2]<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/x9m/exad-implementing-database-machine-x9m.pdf>

Question 2

Question Type: MultipleChoice

Which two statements are true about enabling write-back flash cache?

Options:

- A-** Write-back flash cache is on by default for High Capacity Storage Servers in High Redundancy configurations.
- B-** Write-back Flash Cache cannot be used on Extreme Flash Storage Servers.
- C-** When enabling write-back flash cache in a rolling manner, DBCLI should be used to inactivate the grid disks on all cells first.
- D-** Before write-back flash cache is enabled, you need to drop the Flash Cache first.
- E-** When enabling write-back flash cache in a non-rolling manner, you need to drop the Persistent Memory Commit Accelerator (PMEMLog) first.

Answer:

A, D

Explanation:

[According to the Oracle documentation¹](#), two statements that are true about enabling write-back flash cache are:

Write-back flash cache is on by default for High Capacity Storage Servers in High Redundancy configurations(A). This is because write-back flash cache provides better performance and reliability for data file writes than write-through flash cache.

Before write-back flash cache is enabled, you need to drop the Flash Cache first(D). This is because changing the flash cache mode requires recreating the flash cache with a different configuration.

Question 3

Question Type: MultipleChoice

I/O performance of the prod database on your Exadata Database Machine has degraded slightly over the past month. The database has been allocated to the OLTP I/O Resource Management (IORM) category.

Which two monitoring tools might be useful in examining I/O performance for the prod database?

Options:

- A- OS I/O metrics using Enterprise Manager host pages for the storage servers
- B- OS I/O metrics using OS tools such as iostat on the database servers
- C- I/O-specific dynamic performance views such as v\$iostat_file, v\$iostat_function, and v\$iostat_consumer_group from the prod database instances using SQL *p1us
- D- cellcli (or exaccli/exadcli) to examine storage server metrics such as database, category, ceiidisk, and griddisk
- E- OS I/O metrics using OS tools such as iostat on the storage servers

Answer:

C, D

Explanation:

According to the Oracle documentation¹, two monitoring tools that might be useful in examining I/O performance for the prod database are:

cellcli (or exaclcli/exadcli) to examine storage server metrics such as database, category, ceidisk, and griddisk (D). This tool can help you monitor the I/O Resource Management (IORM) metrics and identify any bottlenecks or imbalances in the storage layer.

I/O-specific dynamic performance views such as v\$iostat_file, v\$iostat_function, and v\$iostat_consumer_group from the prod database instances using SQL *plus. These views can help you monitor the I/O activity and latency at the file, function, and consumer group level.

Question 4

Question Type: MultipleChoice

You are going to create an ACFS Filesystem for use as a staging area for data loads. Examine the following commands:

- 1) `/usr/sbin/mkfs -t acfs /dev/asm/acfs_dataload_vol-123`
- 2) `asmcmd volinfo -G data acfs_dataload_vol | grep Device`
- 3) `acfsutil size +250G -d /dev/asm/acfs_dataload_vol-123 /mnt/acfs_dataload_vol/`
- 4) `sudo /u01/app/19.0.0.0/grid/bin/srvctl start filesystem -device /dev/asm/acfs_dataload_vol-123`
- 5) `sudo chown -R oracle:dba /mnt/acfs_dataload_vol`
- 6) `asmcmd volcreate -G data -s 1024G acfs_dataload_vol`
- 7) `/sbin/acfsutil snap create -w data_load /mnt/acfs_dataload_vol`
- 8) `/u01/app/19.0.0.0/grid/bin/srvctl add filesystem -device /dev/asm/acfs_dataload_vol-123 -path /mnt/a`
`user oracle`
- 9) `sudo /u01/app/19.0.0.0/grid/bin/srvctl stop filesystem -device /dev/asm/acfs_dataload_vol-123`

What steps and in which order should they be executed?

Options:

- A-** 6, 2, 1, 8, 4, 5
- B-** 6, 8, 3, 4, 2, 1
- C-** 6, 3, 1, 9, 7, 8

D- 6, 1, 2, 7, 5, 8

Answer:

A

Explanation:

The steps and the order in which they should be executed to create an ACFS Filesystem for use as a staging area for data loads are:

`asmcrad volcreate -G data -s 1024G acfs_dataload_vol` : This command creates a new volume in the data diskgroup, with a size of 1024 GB, and named `acfs_dataload_vol`.

`asmcmd vol info -G data acfs_dataload_vol | grep Device`: This command will show the device name of the newly created volume

`/usr/sbin/mkfs -t acfs /dev/asm/ac:fs_dataload_vol-123`: This command creates a new ACFS filesystem on the volume created in step 1

`sudo /u01/app/19.0.0.0/grid/bin/srvctl start filesystem -device /dev/asn/acfs_dataload_vol-123`: This command starts the filesystem created in step 3

`sudo chown -R oracle:dba /mnt/acfs_dataload_vol`: This command changes the ownership of the `/mnt/acfs_dataload_vol` to the `oracle:dba` user.

`/sbin/acfsutil snap create -v data_load /mnt/acfs_dataload_vol`: This command creates a snapshot of the filesystem to be used as a staging area.

acfsutil size +2S0G -d /dev/asm/acfs_dataload_vol-123 /mnt/acfs_dataload_vol/ : This command increases the size of the filesystem by 2GB

/u01/app/19-G.O-Q7grid/bin/srvctl add filesystem -device /dev/asm/acfs_dataload_vol-

Question 5

Question Type: MultipleChoice

You are adding a disk expansion kit to a running Exadata X8M Database Machine's Database Servers, and have a filesystem layout that includes:

Filesystem	Mounted on
/dev/mapper/VGExaDb-LVDbSys1	/
/dev/mapper/VGExaDb-LVDbVar1	/var
/dev/mapper/VGExaDb-LVDbHome	/home
/dev/mapper/VGExaDb-LVDbTmp	/tmp
/dev/mapper/VGExaDb-LVDbVarLog	/var/log
/dev/mapper/VGExaDb-LVDbOra1	/u01
/dev/mapper/VGExaDb-LVDbVarLogAudit	/var/log/audit

After running the following commands, which command needs to be run to add 20G of space to the filesystem mounted on /u01?

```
# parted -s /dev/sda mkpart primary 240132160s 8189439966s
# parted -s /dev/sda set 3 lvm on
# lvm pvcreate --force /dev/sda3
# lvm vgextend VGExaDb /dev/sda3
```

Options:

A- # lvextend -L +20G ---verbose /dev/mapper/VGExaDb-LVDbOral

B- # xfs_growfs /u01 +20G

C- # resize2fs +20G /dev/VGExaDb/LVDbOral

D- # lvextend -L +20G ---verbose /dev/VGExaDb/LVDbOral

Answer:

A

Explanation:

After running the commands above, the filesystem mounted on /u01 is on the logical volume /dev/mapper/VGExaDb-LVDbOral. So, to add 20G of space to the filesystem mounted on /u01, the command that needs to be run is:

```
lvextend -L +20G --verbose /dev/mapper/VGExaDb-LVDbOral
```

This command will extend the logical volume /dev/mapper/VGExaDb-LVDbOral by 20 GB of space. It is important to note that the option --verbose is used to display the progress of the operation.

Question 6

Question Type: MultipleChoice

Which three steps are required to expand an Exadata X9M-2 Database Server with the lowest memory configuration available to the highest memory configuration?

Options:

- A- Add 12x 32GB DIMMs.
- B- Add 12x 64GB DIMMs.
- C- Shutdown the Database Server if running.
- D- Add 16x 32GB DIMMs.
- E- Add 16x 64GB DIMMs.
- F- Memory cannot be expanded on Exadata X9M-2 Database Servers.

G- Remove existing memory modules.

H- Add 32x 64GB DIMMs.

I- Add 24x 32GB DIMMs.

Answer:

C, G, H

Explanation:

<https://www.oracle.com/a/ocom/docs/engineered-systems/exadata/exadata-x9m-2-ds.pdf>

<https://chriscraftoracle.wordpress.com/2022/11/09/what-makes-exadata-faster/>

Question 7

Question Type: MultipleChoice

Which are two correct statements for managing virtual deployment using Oracle Exadata Deployment Assistant (OEDA)?

Options:

- A- OEDA allows customers to have both bare metal (BM) and virtual machine (VM) in an Exadata X9M Quarter Rack.
- B- There is no limit on the number of VMs in an Exadata rack as long as the Exadata rack has adequate resources.
- C- OEDA deployment steps include calibrate cells, create cell disks, and resecure machine.
- D- OEDA sets up key-based authentication for the root user by using the setuprootssh.sh utility included with OEDA.

Answer:

C, D

Explanation:

C)OEDA deployment steps include calibrate cells, create cell disks, and resecure machine2.

D)OEDA sets up key-based authentication for the root user by using the setuprootssh.sh utility included with OEDA2.

Question 8

Question Type: MultipleChoice

Which dbmcli command is NOT valid on Exadata X9M?

Options:

- A- dbmcli -e 'LIST METRICHISTORY WHERE name LIKE 'DS_.*''
- B- dbmcli -e 'LIST METRICCURRENT WHERE name = 'DS_TEMP' "
- C- dbmcli -e 'LIST IBPORT DETAIL'
- D- dbmcli -e 'LIST ALERTHISTORY WHERE ageInMinutes < 15'

Answer:

C

Explanation:

The dbmcli -e 'LIST IBPORT DETAIL' command is not valid on Exadata X9M. According to the Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book[1][2], this command is only valid on Exadata X3 and X4 models. The other three commands are valid on Exadata X9M.

<https://forums.oracle.com/ords/apexds/map/tech/apps-infra/discussion/4342715/monitoring-roce-performance>

Question 9

Question Type: MultipleChoice

A new Exadata Quarter Rack with 2 Database Servers and 3 HC Storage Servers and 3-phase 15kVA PDUs is being installed in a Data Center. However, the Data Center is only providing enough power for a single cable from each PDU.

Which statement is correct?

Options:

- A-** The installation can go ahead, no change is required.
- B-** The installation cannot proceed until two power feeds are available per PDU.
- C-** A splitter cable can be used to provide power to all PDU cables.
- D-** The power cables from the servers to the PDUs can be rearranged inside the rack following OECA guidance to utilize a single PDU power cable.

Answer:

A

Explanation:

According to the Oracle Exadata Database Machine X9M-2 Data Sheet¹, each PDU has two power input cables that can be connected to separate power sources for redundancy. However, if only one power source is available per PDU, then a single cable can be used to provide sufficient power for normal operation.

<https://www.oracle.com/a/ocom/docs/engineered-systems/exadata/exadata-x9m-2-ds.pdf>

Question 10

Question Type: MultipleChoice

Which two sections of the AWR report shows statistics for X9M Persistent Memory Cache?

Options:

- A- PMEM Pool cache Read Hits in the Cache Sizes portion of the Report Summary
- B- PMEM Pool Misses in the Exadata Outlier Summary
- C- cell PMEM cache Read Hits in the Database IOs portion of the Performance Summary
- D- PMEM Cache section within Memory Statistics
- E- PMEM Cache section within Exadata Smart Statistics

Answer:

C, E

Explanation:

Option C shows the number of read hits from PMEM cache on storage servers which indicates how much data was served from PMEM instead of flash or disk1.

Option E shows detailed information about PMEM cache such as size, utilization, hit ratio, read latency and write latency2.

Question 11

Question Type: MultipleChoice

Examine these commands:

1. Execute "crscti stop cluster -all" as the grid user from one database server.
2. Execute "crscti stop cluster -all" as root from one database server.
3. Power off all network switches.
4. Execute "crscti stop cluster" as root from one database server.

5. Execute "crscti stop cluster" as the grid user from one database server.
6. Power off the rack using the power switches on the PDUs.
7. Execute "shutdown -h now" on all database servers.
8. Execute "shutdown -h now" on all Exadata storage servers.

Which is the correct order or the required commands to completely power off an Exadata Database Machine in an orderly fashion?

Options:

- A-** 5, 8, 7, and 6
- B-** 4, 7, 8, 3, and 6
- C-** 2, 8, 7, 3, and 6
- D-** 2, 7, 8, and 6
- E-** 1, 8, 7, 3 and 6

Answer:

A

Explanation:

Option 5 stops the clusterware on all database servers as the grid user3.

Option 8 shuts down all Exadata storage servers using their ILOM interfaces1.

Option 7 shuts down all database servers using their ILOM interfaces1.

Option 6 powers off the rack using the power switches on the PDUs2.

Question 12

Question Type: MultipleChoice

Which of the following is NOT a requirement when validating, receiving, unpacking, and planning access route and space requirements for Exadata Database Machine?

Options:

- A-** The entire access route to the installation site should be free of raised-pattern flooring that can cause vibration.
- B-** 914mm of space required above the rack height is required for maintenance access.
- C-** The incline of any access route ramp must be less than or equal to 6 degrees.

- D-** All four leveling and stabilizing feet should be raised and out of the way prior to moving the rack.
- E-** Oracle Exadata Rack may only be installed on raised floor environments.
- F-** A conditioned space is required to remove the packaging material to reduce particles before entering the data center.

Answer:

E

Explanation:

Exadata Database Machine is a pre-configured and pre-tuned hardware and software system designed to run Oracle Database, it can be installed on raised floor environments, but also on concrete or tile floors

Oracle Exadata Database Machine X9M Implementation Essentials states that Exadata racks are designed to be installed on flat surfaces, not raised floor environments. It is not required to install the rack on raised floor environments. Additionally, the other options listed are all requirements for validating, receiving, unpacking, and planning access route and space requirements for Exadata Database Machine. (Source: Oracle Exadata Database Machine X9M Implementation Essentials, page 41)

<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmin/index.html>

To Get Premium Files for 1Z0-902 Visit

<https://www.p2pexams.com/products/1z0-902>

For More Free Questions Visit

<https://www.p2pexams.com/oracle/pdf/1z0-902>

