



# 1Z0-902<sup>Q&As</sup>

Oracle Exadata Database Machine X9M Implementation Essentials

**Pass Oracle 1Z0-902 Exam with 100% Guarantee**

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

<https://www.passapply.com/1z0-902.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

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

- 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

Correct Answer: CE

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 2

For which four component failures on an X9M Database Machine does Auto Service Request (ASR) raise service requests?

- A. RoCE network interface cards in the storage servers
- B. fans in the storage servers
- C. Cisco RDMA over Converged Ethernet (RoCE) switches
- D. RoCE network interface cards in the database servers
- E. power distribution units
- F. Cisco management switch
- G. power supplies in the database servers

Correct Answer: ACEG

Explanation: According to the Oracle Auto Service Request (ASR) documentation1, ASR raises service requests for qualified Oracle products that are detected with specific faults. The qualified Exadata products include2:

Database servers

Storage servers

InfiniBand switches

Cisco switches (X8M and later systems)

Power distribution units (PDUs)



### QUESTION 3

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?

- 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 exacl/exadcli) to examine storage server metrics such as database, category, ceidisk, and griddisk
- E. OS I/O metrics using OS tools such as iostat on the storage servers

Correct Answer: CD

Explanation: According to the Oracle documentation<sup>1</sup>, two monitoring tools that might be useful in examining I/O performance for the prod database are:

cellcli (or exacl/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 \*p1us ? These views can help you monitor the I/O activity and latency at the file,

function, and consumer group level.

---

### QUESTION 4

You are in the process of upgrading your nonvirtualized X9M-2 Database Machine elastic configuration with 4 database servers and 7 HC storage servers with an additional 4 database servers and 7 HC storage servers.

The new storage servers are called DM01CEL08 through dmoicel14.

After creating 96 new griddisks, you issued this SQL statement:



```
SQL> ALTER DISKGROUP DATA ADD DISK
 2> 'O/*/DATA*DM01CEL08*'
 3> 'O/*/DATA*DM01CEL09*'
 4> 'O/*/DATA*DM01CEL10*'
 5> 'O/*/DATA*DM01CEL11*'
 6> 'O/*/DATA*DM01CEL12*'
 7> 'O/*/DATA*DM01CEL13*'
 8> 'O/*/DATA*DM01CEL14*'
 9> REBALANCE POWER 512;
```

How many failgroups if any, will be added to the DATE diskgroup by executing this SQL statement?

- A. 1 consisting of all 96 griddisks
- B. 96 consisting of one griddisk each
- C. 0 because the new griddisks will be added to the existing failgroups
- D. 12 consisting of seven griddisks each
- E. 7 consisting of 12 griddisks each

Correct Answer: A

Explanation: This SQL statement is adding the new griddisks to the existing diskgroup "DATA" and creating one new failgroup, consisting of all 96 griddisks. The "REBALANCE POWER 512" option tells the system to perform a rebalance operation with a power of 512. It means the system distributes the data evenly across all the disks in the diskgroup using a power of 512.

## QUESTION 5

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
```

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

Correct 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.

[1Z0-902 VCE Dumps](#)

[1Z0-902 Practice Test](#)

[1Z0-902 Study Guide](#)