PassQueen
a.k.a dumpsqueen.com

Oracle Exadata Database Machine X9M Implementation Essentials

Oracle 1z0-902

Version Demo

Total Demo Questions: 10

Total Premium Questions: 106

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Exadata Architecture	18
Topic 2, Exadata Installation	15
Topic 3, Exadata Configuration	18
Topic 4, Exadata Administration	45
Topic 5, Exadata Security	5
Topic 6, Exadata Patching	5
Total	106

QUESTION NO: 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

ANSWER: C E

Explanation:

The AWR report exposes X9M Persistent Memory (PMEM) Cache activity in both the high-level database I/O summary and the Exadata-specific smart I/O statistics. **cell PMEM cache Read Hits in the Database IOs portion of the Performance Summary** is the database-facing measure of physical reads served from the cells' PMEM cache. It helps quantify how much read activity avoids slower storage media by being fulfilled through persistent memory.

PMEM Cache section within Exadata Smart Statistics is also correct because Exadata-specific AWR reporting groups smart-storage and cache metrics in that area, including PMEM cache activity. This information is useful when evaluating the effectiveness of Exadata offload and cache tiers, and when correlating database I/O behavior with cell-side performance. Oracle describes AWR as the repository and reporting mechanism for database workload and performance statistics, while Exadata documentation covers monitoring the database machine and its cell-side performance metrics. See Oracle's Managing Automatic Workload Repository documentation and the Exadata monitoring documentation.

QUESTION NO: 2

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

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:

"Oracle Exadata Rack may only be installed on raised floor environments." is the item that is not a requirement. Exadata Database Machine racks can be deployed in either raised-floor or non-raised-floor data centers, provided the selected site meets the applicable floor-loading, clearance, environmental, power, cooling, and service-access specifications. A raised floor is therefore not an exclusive installation prerequisite. Installation planning must instead validate that the floor can safely support the fully populated rack and that the room layout permits delivery, positioning, cabling, airflow, and future maintenance. Oracle's installation documentation treats raised-floor and non-raised-floor deployments as supported site configurations, with planning considerations appropriate to each type of floor. The key implementation principle is to assess the actual site against Oracle's physical installation requirements, rather than require one particular flooring design. This makes the statement's "may only" restriction incorrect as a requirement and identifies it as the correct response to the question. See Oracle's Exadata Database Machine documentation library and the site preparation guidance for the current physical installation and facility-planning requirements.

QUESTION NO: 3

Which two statements are true about the IPTables firewall configuration on a new Exadata Database Machine after initial deployment?

A. IPTables is installed and available but not configured or running on the storage servers.

B. IPTables is configured and running with Oracle-supplied rules on the storage servers.

C. IPTables is configured and running with Oracle-supplied rules on the database servers.

D. IPTables is installed and available but not configured or running on the database servers.

E. IPTables is installed and available but not configured on any server.

ANSWER: A C

Explanation:

After initial deployment, Exadata distinguishes between database servers and storage servers for host-firewall configuration. **IPTables is installed and available but not configured or running on the storage servers.** This leaves the storage-server firewall available for a customer to configure when its security architecture requires host-based filtering, while avoiding an out-of-the-box rule set that could affect the specialized storage and interconnect communications required by the cells.

IPTables is configured and running with Oracle-supplied rules on the database servers. Oracle supplies and enables the initial database-server rule set as part of the deployed Exadata environment. These rules provide baseline host protection while allowing the network services and cluster/database communications needed for normal database-server operation. Administrators should preserve the required Exadata connectivity when reviewing or extending firewall policy, and should use Oracle-supported procedures rather than replacing the deployment configuration indiscriminately.

This default split is important operationally: firewall status must be assessed separately on database and storage servers rather than assumed to be identical across all Exadata hosts. Oracle's Exadata documentation is available at Oracle Exadata Database Machine documentation; the underlying Oracle Linux IPTables administration concepts are documented in Oracle Linux Security Guide.

QUESTION NO: 4

Which two of the following network connection types can be deployed on the client network ports of an Exadata X9M-2 Database Server?

A. 2x 25Gb ports on eth1 and eth2 using SFP28 transceivers and OM4 fiber cable for backup network, and 2x 25Gb ports on eth5 and eth6 using
SFP28 transceivers and OM4 fiber cable for client network

B. 2x 25Gb ports on eth1 and eth2 using SFP28 transceivers and OM4 fibre cable for client network, and 2x 10Gb ports on eth9 and eth10 using RJ45 cat6 cable for backup network

C. 2x 10Gb ports on eth1 and eth2 using FU45 cat6 cable for client network, and 2x 25Gb ports on eth11 and eth12 using SFP28 transceivers and OM4 fiber cable for backup network

D. 2x 25Gb ports on eth1 and eth2 using SFP28 transceivers and OM4 fiber cable for backup network, and 2x 10Gb ports on eth3 and eth4 using RJ45 cat6 cable for client network

ANSWER: A B

Explanation:

The supported client-network deployments described are **2x 25Gb ports on eth1 and eth2 using SFP28 transceivers and OM4 fibre cable for client network, and 2x 10Gb ports on eth9 and eth10 using RJ45 cat6 cable for backup network** and **2x 25Gb ports on eth1 and eth2 using SFP28 transceivers and OM4 fiber cable for backup network, and 2x 25Gb ports on eth5 and eth6 using SFP28 transceivers and OM4 fiber cable for client network**. The X9M-2 database server provides 25GbE SFP28 interfaces that can be assigned in supported pairs to the client and backup networks. Therefore, the

client network can use the eth1/eth2 25GbE pair while the backup network uses the separate 10GbE RJ45 pair, or it can use the eth5/eth6 25GbE pair while eth1/eth2 serves the backup network. In each case, using a pair provides the redundant network connectivity expected for Exadata client access. SFP28 optical transceivers with OM4 fiber are appropriate for the 25GbE optical connections, while the 10GbE Base-T backup connection uses RJ45 Category 6 cabling. Oracle's X9M-2 database-server specification and Exadata deployment documentation define these supported physical network interfaces and cabling choices. See the Exadata X9M-2 Database Server documentation and Oracle Exadata documentation.

QUESTION NO: 5

Which three statements are true about an ASM disk group rebalance operation?

A. ASM automatically initiates a rebalance after an ASM disk is added or dropped.

B. The progress of an active rebalance can be monitored using V$ASM_OPERATION.

C. The rebalance power setting affects the resources allocated to the rebalance operation.

D. A disk group must be dismounted before ASM can rebalance it.

E. ASM rebalancing is performed only after the database instances are shut down.

F. A rebalance operation permanently disables ASM mirroring.

ANSWER: A B C

Explanation:

ASM automatically initiates a rebalance after an ASM disk is added or dropped is correct. ASM redistributes extents to reflect the changed disk-group membership and maintain balanced allocation across available disks.

The progress of an active rebalance can be monitored using V$ASM_OPERATION is correct. This dynamic performance view reports information about active ASM operations, including rebalance progress and estimated work.

The rebalance power setting affects the resources allocated to the rebalance operation is correct. Higher power values can complete the operation more quickly but can also increase the I/O and CPU resources consumed by ASM.

See Maintaining Oracle ASM Disk Groups.

QUESTION NO: 6

Which ASM disk group attribute controls the minimum Oracle ASM software version required to mount an ASM disk group?

A. COMPATIBLE.ASM

B. COMPATIBLE.RDBMS

C. ASM_POWER_LIMIT

D. AU_SIZE

E. SECTOR_SIZE

ANSWER: A

Explanation:

COMPATIBLE.ASM specifies the minimum Oracle ASM version required to mount an ASM disk group. It also controls which ASM disk-group features can be enabled. Increasing this attribute can enable newer ASM functionality, but the value cannot be reduced afterward.

The related `COMPATIBLE.RDBMS` attribute controls the minimum database software version that can use the disk group; it does not determine the minimum ASM version required to mount it. Oracle documents these attributes in the Oracle ASM Disk Group Compatibility documentation.

QUESTION NO: 7

Examine this list of software components:

1. Oracle KVM Guest

2. Oracle Enterprise Manager Agent (OMA)

3. ASM instance

4. RDBMS instance

5. Automatic Diagnostic Repository Command Interpreter (ADRCI)

6. CELLCLI

7. Cell Server(CELLSRV)

8. diskmon

9. Restart Server (RS)

10. Management Server (MS)

What is the correct location where these software components can run in the standard Exadata Database Machine deployment?

A. 2, 3, 4, 8, and 10 run on the database servers; 1, 5, 6, 7and 9 run on the Exadata storage servers.

B. 1, 2, 3, 4, 9 and 10 run on the database servers; 5, 6, 7, 8, 9, and 10 run on the Exadata storage servers.

C. 1, 2, 3, 4, 5, 8, 9 and 10 run on the database servers; 5, 6,7, 9 and 10 run on the Exadata storage servers.

D. 3, 4, 8, and 10 run on the database servers; 1, 2, 5, 6, 7 and 9 run on the Exadata storage servers.

E. 1, 2, 3, 4, 8 and 9 run on the database servers; 5, 6, 7, 9 and 10 run on the Exadata storage servers.

ANSWER: C

Explanation:

The correct placement is that Oracle KVM Guest, Oracle Enterprise Manager Agent, ASM instance, RDBMS instance, ADRCI, diskmon, Restart Server, and Management Server can run on database servers; ADRCI, CELLCLI, Cell Server, Restart Server, and Management Server can run on Exadata storage servers. Database servers host the database software stack, including the Oracle Database and ASM instances, along with diskmon, which supports monitoring of Exadata storage. A KVM guest is also a database-server workload in an Exadata deployment that uses virtualization.

Exadata storage servers run the Exadata System Software stack. CELLSRV is the core cell service that provides intelligent storage processing, while CellCLI is the command-line administration interface for managing cell resources. The management and restart services support management operations and service availability, and ADRCI is available where diagnostic data must be examined. This division reflects Exadata's architecture: database servers execute database and ASM processing, while storage servers provide specialized cell services and storage offload capabilities. Oracle documents the database-server and storage-server software roles in the Oracle Exadata Database Machine documentation and the Oracle Exadata System Software documentation.

QUESTION NO: 8

You have been asked to design a backup solution for an Exadata X9M-2 Quarter Rack with Extreme Flash Storage Servers connected to a new ZFS Storage Appliance ZS7 with 2 Storage Controllers with 100Gb Ethernet cards and 3 Storage Trays. You are using Oracle Exadata Configuration Assistant to validate the rack layout.

1. Use "Add Equipment" to add the Exadata X9M EF Storage Servers, starting from RU10.

2. Use drop down to add ZFS Storage Appliance Controllers.

3. You cannot add ZFS Storage Appliance to an Exadata Rack.

4. Use "Add Equipment" to add the ZFS Storage Trays, starting from RU1.

5. Use drop down to add ZFS Storage Trays.

6. Use "Add Equipment" to add the Exadata X9M Database Servers, starting from RU16.

7. Use "Add Equipment" to add the Exadata X9M EF Storage Servers, starting from RU1.

8. Use drop down to add Exadata X9M EF Storage Servers.

9. Use "Add Equipment" to add the ZFS Storage Controllers, starting from RU27.

10. Use drop down to add Exadata X9M Database Servers.

11. Use "Add Equipment" to add the ZFS Storage Trays, starting from RU31.

Which of these steps are correct and what is their correct order?

A. 10, 8, 2, 5

B. 4,1, 6, 9

C. 3

D. 10, 8, 9, 11

E. 7,6,9,11

ANSWER: D

Explanation:

The correct sequence is **10, 8, 9, 11**. Oracle Exadata Configuration Assistant is used to construct and validate the physical rack layout by selecting the appropriate Exadata system components and then placing additional equipment at its required rack-unit positions. For this X9M-2 Quarter Rack design, first select the Exadata X9M Database Servers from the drop-down list. Next, select the Exadata X9M Extreme Flash Storage Servers from the drop-down list. These selections establish the native Exadata server configuration before the external backup appliance equipment is placed in the rack model.

After the Exadata components have been selected, use *Add Equipment* to place the two ZFS Storage Appliance controllers beginning at RU27. Then use *Add Equipment* to place the three ZFS Storage Trays beginning at RU31. This order reflects the intended physical installation arrangement: the ZFS controllers occupy the rack units above the storage trays, while the trays are installed in the lower portion of the rack. Modeling the equipment this way lets Configuration Assistant validate rack space and the planned physical layout for the ZFS-based backup solution. Oracle describes Exadata configuration and related deployment tasks in its Exadata documentation; ZFS appliance hardware and deployment information is available in the Oracle ZFS Storage Appliance documentation.

QUESTION NO: 9

Which are two valid reasons for executing an X9M-2 Exadata storage server rescue procedure?

A. the failure of physical disk 1

B. corruption in the / (root) filesystem

C. the failure of both physical M.2 disks

D. the failure of physical disk 0 and 11

E. moving all disks from one cell to another as part of a chassis-level component failure

F. accidental loss of all data from all griddisks in a storage server

G. corruption in a normal or high redundancy ASM diskgroup

ANSWER: B C

Explanation:

Corruption in the / (root) filesystem is a valid reason to perform an X9M-2 storage server rescue because the operating-system filesystem is essential to booting and operating the cell. The rescue process supplies a controlled method to restore the storage server software environment when that local system installation cannot be reliably repaired through ordinary maintenance actions. It is therefore focused on recovering the cell's boot and operating-system capability, rather than rebuilding database or ASM data.

The failure of both physical M.2 disks is also a valid reason. On an X9M-2 storage server, the M.2 devices provide the mirrored boot storage for the operating system. Loss of both devices removes the local boot media and requires the rescue workflow to restore the server to a usable, supported state. Oracle's Exadata maintenance documentation describes rescue procedures as recovery mechanisms for serious local software or boot-device failures and directs administrators to follow the platform-specific procedure and subsequent validation steps. See the Oracle Exadata Database Machine Maintenance Guide and Oracle Exadata Database Machine documentation.

QUESTION NO: 10

Which ASM initialization parameter controls the power level used by Automatic Storage Management rebalance operations?

A. ASM_POWER_LIMIT

B. ASM_DISKSTRING

C. ASM_PREFERRED_READ_FAILURE_GROUPS

D. DISK_REPAIR_TIME

ANSWER: A

Explanation:

`ASM_POWER_LIMIT` specifies the default power level for ASM rebalancing operations. A higher value permits ASM to use more resources to complete a rebalance more quickly, while a lower value reduces the impact of rebalancing on database I/O. The setting can be overridden for an individual rebalance operation by specifying a POWER clause in the relevant ASM command.

ASM rebalancing occurs when disks are added, dropped, resized, or otherwise changed in an ASM disk group. See Oracle ASM disk group maintenance documentation.