

Oracle 1Z0-902

Oracle Exadata Database Machine X9M Implementation Essentials

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

1. Up to Date products, reliable and verified.
2. Questions and Answers in PDF Format.



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/1z0-902>

Latest Version: 7.0

Question: 1

What is the maximum DRAM capacity you can expand an X9M-2 DB Server?

- A. 1536GB
- B. 512GB
- C. 2048GB
- D. 768GB
- E. 1024GB
- F. 384GB

Answer: A

Explanation:

According to 1, Exadata X9M-2 Database Servers are based on Intel Xeon Platinum 8368Y processors with 32 cores and 768 GB of DRAM memory per server. This memory can be expanded up to 1.5 TB per server by adding 24 x 32 GB DIMMs. The DRAM memory is used for buffer cache, PGA, and other database purposes¹. Exadata X9M-2 Database Servers also have 25.6 TB of Persistent Memory Acceleration (PMEM) per server, which is used for log writes and columnar data caching¹. PMEM is a new type of memory that combines the speed of DRAM with the persistence of flash². Exadata X9M-2 Database Servers are designed to run Oracle Database workloads with high performance, scalability, and reliability²

Question: 2

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, 7 and 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:

Oracle KVM Guest: This is a virtual machine that runs on top of Oracle Linux KVM hypervisor. It can be used to run Oracle Database or other applications on Exadata Database Machine2. Therefore, it runs on the Database Servers.

Oracle Enterprise Manager Agent (OMA): This is a software agent that communicates with Oracle Enterprise Manager Cloud Control and provides monitoring and management capabilities for Exadata Database Machine2. Therefore, it runs on both Database Servers and Storage Servers.

ASM instance: This is an instance of Oracle Automatic Storage Management (ASM), which is a volume manager and a file system for Oracle Database files. It manages disk groups that span across multiple Storage Servers2. Therefore, it runs on the Database Servers.

RDBMS instance: This is an instance of Oracle Database that processes SQL statements and executes transactions. It uses ASM disk groups to store data files, control files, redo log files, etc2. Therefore, it runs on the Database Servers.

Automatic Diagnostic Repository Command Interpreter (ADRCI): This is a command-line tool that enables you to view diagnostic data stored in the Automatic Diagnostic Repository (ADR). ADR is a filebased repository for database diagnostic data such as trace files, alert logs, etc2. Therefore, ADRCI runs on both Database Servers and Storage Servers, depending on where the ADR is located.

CELLCLI: This is a command-line interface that enables you to configure and manage Exadata Storage Server Software. It allows you to perform tasks such as creating disk groups, monitoring cell health, applying patches, etc2. Therefore, it runs on the Storage Servers.

Cell Server(CELLSRV): This is a process that runs on each Storage Server and handles I/O requests from the Database Servers. It implements Exadata Smart Scan , which offloads data-intensive SQL operations from the Database Servers to the Storage Servers2. Therefore, it runs on the Storage Servers.

diskmon: This is a process that monitors the status of disks and flash devices on each Storage Server. It reports disk failures and performs automatic disk reclamation2. Therefore, it runs on the Storage Servers.

Restart Server (RS): This is a process that manages automatic restarts of critical processes such as CELLSRV , MS , or OMA in case of failures. It also handles graceful shutdowns and startups of all processes on each server2. Therefore, it runs on both Database Servers and Storage Servers.

Management Server (MS): This is a process that provides management services for each server such as collecting metrics , logging events , executing commands from CELLCLI , etc2. Therefore, it runs on both Database Servers and Storage Servers

Question: 3

Your customer needs to ensure that their data is available on the Exadata machine during updates. The customer wants to be able to update one server at a time but still be protected against single-node

server failure.

What ASM redundancy level should they use?

- A. Normal
- B. Sparse
- C. High
- D. External
- E. Extended

Answer: A

Explanation:

Your customer needs to ensure that their data is available on the Exadata machine during updates. The customer wants to be able to update one server at a time but still be protected against single-node server failure. This means that they need a redundancy level that can tolerate at least one disk failure without losing data availability or performance. Therefore, the best option for them is normal redundancy . Normal redundancy provides two-way mirroring by default, which means that each data extent has a primary copy and a mirror copy on different disks in different failure groups¹. A failure group is a subset of disks within a disk group that are likely to fail at the same time due to common dependencies². On Exadata Database Machine , each Storage Server is considered as a failure group by default². This means that if one Storage Server fails or needs to be updated , the data on its disks can still be accessed from another Storage Server that has the mirror copies of the same data extents². Therefore, normal redundancy ensures data availability and performance during updates or single-node server failures.

<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/sagug/exadataadministering-asm.html>

Question: 4

Which three statements are true about the CELLCLI command?

- A. It can execute commands on multiple storage servers in parallel.
- B. It has command-line history.
- C. It requires root privileges to create CELLDISKS and GRIDDISKS.
- D. It can be executed on storage servers.
- E. It can be executed using the DCLI utility.

Answer: BDE

Explanation:

The statements that are true about the CellCLI command are:

It has command-line history . According to 1, you can use the up and down arrow keys to scroll through previous commands that you entered in CellCLI . You can also use Ctrl+R to search for a previous command by entering part of it .

It can be executed on storage servers . According to 1, you can invoke CellCLI by logging in to the Exadata cell as cellmonitor , celladmin , or root user and typing “cellcli” . Within CellCLI , you can perform various management and monitoring tasks for the cell134.

It can be executed using the DCLI utility . According to 5, DCLI (distributed CLI) is another utility that can be used to execute scripts and commands across multiple storage cells from a single interface. You can use DCLI to run CellCLI commands on multiple cells at once54.

Question: 5

An Exadata X9M-2 Elastic Rack with 4 Database Servers and 8 HC Storage Servers and 3-phase 24kVA 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?

- A. A splitter cable can be used to provide power to all PDU cables.
- B. 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.
- C. The installation cannot proceed until two power feeds are available per PDU.
- D. The installation can go ahead, no change is required.

Answer: C

Explanation:

In order for an Exadata X9M-2 Elastic Rack to be installed, two power feeds are required for each of the three-phase 24kVA PDUs. A single cable from each PDU will not be enough to power the rack, and neither a splitter cable nor rearrangement of the power cables from the servers to the PDUs will be able to provide sufficient power for the system. Therefore, the installation cannot proceed until two power feeds are available per PDU.

For more information about power requirements for Exadata X9M-2 Elastic Racks, refer to the Oracle Exadata Database Machine X9M Implementation Essentials official text book and resources.

Search results: [1] Oracle Exadata Database Machine X9M-2 Full Rack Installation - Oracle Docs [2] Exadata Database Machine X8M-2 Mid-Size Rack Installation - Oracle Docs

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

Thank You for Trying Our Product
Special 16 USD Discount Coupon: NSZUBG3X
Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>