

1Z0-070^{Q&As}

Oracle Exadata X5 Administration

Pass Oracle 1Z0-070 Exam with 100% Guarantee

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

<https://www.pass2lead.com/1z0-070.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

An Exadata storage server physical disk on an X5 high-capacity full rack entered the predictive failure state.

Identify the two steps that you must perform to replace this failed physical disk.

- A. Add the griddisks back into the ASM diskgroup they used to be a member of.
- B. Create a new celldisk and new griddiskson the replaced physical disk.
- C. Verify that the griddisks located on the physical disk have been successfully dropped from the associated ASM diskgroups.
- D. Identify the griddisks located on the failed physical disk and drop them from the associated ASM diskgroups.
- E. Replace the failed physical disk.

Correct Answer: CE

Explanation: You may need to replace a physical disk because the disk is in warning - predictive failure status. The predictive failure status indicates that the physical disk will soon fail, and should be replaced at the earliest opportunity. The Oracle ASM disks associated with the grid disks on the physical drive are automatically dropped, and an Oracle ASM rebalance relocates the data from the predictively failed disk to other disks.

Note: After the physical disk is replaced, the grid disks and cell disks that existed on the previous disk in that slot are re-created on the new physical disk. If those grid disks were part of an Oracle ASM group, then they are added back to the disk group, and the data is rebalanced on them, based on the disk group redundancy and ASM_POWER_LIMIT parameter.

References: http://docs.oracle.com/cd/E80920_01/DBMMN/maintaining-exadata-storageservers.htm#DBMMN21047

QUESTION 2

Which two options should be used to identify a damaged or failing flash card on an X6 Database Machine high-capacity storage server?

- A. Using the CELLCLI CALIBRATE command on the storage server after logging in as the root user
- B. Using the CELLCLI LIST LUN DETAIL command as the celladmin user
- C. Using the CELLCLI CALIBRATE command on the storage server after logging in as the celladmin user
- D. Hardware monitoring using the storage server ILOM

Correct Answer: AD

Explanation:

The CALIBRATE command runs raw performance tests on cell disks, enabling you to verify the disk performance before the cell is put online. You must be logged on to the cell as the root user to run CALIBRATE.

Oracle ILOM automatically detects system hardware faults and environmental conditions on the server. If a problem occurs on the server, Oracle ILOM will automatically do the following:

Illuminate the Service Required status indicator (LED) on the server front and back panels.

Identify the faulted component in the Open Problems table.

Record system information about the faulted component or condition in the event log.

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-cellcli.htm#SAGUG20617

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-cellcli.htm#SAGUG20617

QUESTION 3

You are planning your deployment of Enterprise Manager to monitor all the components of an X5 Database Machine.

A part of the requirement is to provide for high availability of the monitoring infrastructure.

If the host running the agent that has Database Machine targets bound to it fails, the monitoring of these targets must be done by another agent.

Which three statements are true regarding the configuration used to support this requirement?

- A. Database Machine plug-ins must be deployed to at least two Enterprise Manager agents.
- B. Fail back to the original agent when the host is restarted is done automatically.
- C. Fail over to any secondary agent is done automatically.
- D. Fail over to any secondary agent must be done manually.
- E. Database machine plug-ins must be deployed to all Enterprise Manager agents.
- F. Fail back to the original agent when the host is restarted must be done manually.

Correct Answer: BCE

QUESTION 4

You issued these commands to all Exadata Storage Servers in an X6 Exadata Database Machine using dcli:

```
alter iormplan objective = offalter iormplan active
```

There are no database or category plans defined.

You are suffering I/O performance problems at certain times, which vary by day and week.

DSS and batch workloads perform well, but your OLTP workloads suffer poor response times when running at the same time as these other workloads.

You must ensure that DSS and batch workloads retain relatively high throughput without causing excessive degradation of OLTP performance.

Which statement would you issue to all Exadata Storage Servers to achieve this?

- A. alter iormplan objective = low_latency
- B. alter iormplan objective = high_throughput
- C. alter iormplan objective = balanced
- D. alter iormplan objective = auto

Correct Answer: A

Explanation:

When the objective option is set to basic, the database resource plan maximum utilization limits are not enforced. For stricter plan conformance, and enforcement of maximum utilization limits, the objective option must be set to something other than basic. The supported IORM objectives are auto, low_latency, balanced, and high_throughput.

The objective may also be changed to low_latency in which case the latency for critical I/Os is very good but there is a significant degradation in scan throughput when both workloads are running concurrently.

Incorrect Answers:

B: The objective can be changed to high_throughput which will increase the scan throughput at the cost of critical I/O latency.

C, D: The recommended objective option is auto which allows IORM to continuously monitor the workloads, and select the best mode based on the active workloads currently on the cells. The objective values of auto and balanced have the same behavior.

QUESTION 5

You applied patches to both the storage servers and the database servers in your Exadata X6 Database Machine.

You use Database Replay for a 90-minute workload to verify that there are no performance regressions after applying the patches.

Which three options could you use to immediately measure the performance and to check for possible regressions caused by changes in the storage server software?

- A. cellcli list metriccurrent

- B. SQL Access Advisor
- C. cellcli list alerthistory
- D. SQL trace using dbms_monitor
- E. cellcli list metrhistory

Correct Answer: ABD

Reference:

<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/sagug/exadata-storage-server-cellcli.html#GUID-9003D6A0-2113-4846-9471-572A0CFBAD9E>

<https://docs.oracle.com/en/database/oracle/oracle-database/19/tgsql/introduction-to-sql-tuning.html#GUID-B237A2A2-7CA2-4AC9-B81C-71C5757C20B2>

<https://docs.oracle.com/en/database/oracle/oracle-database/19/tgsql/performing-application-tracing.html#GUID-33FA4E54-C759-480B-A075-CF75E6D85CAD>

<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/sagug/exadata-storage-server-cellcli.html#GUID-9003D6A0-2113-4846-9471-572A0CFBAD9E> <https://docs.oracle.com/en/database/oracle/oracle-database/19/tgsql/introduction-to-sql-tuning.html#GUID-B237A2A2-7CA2-4AC9-B81C-71C5757C20B2> <https://docs.oracle.com/en/database/oracle/oracle-database/19/tgsql/performing-application-tracing.html#GUID-33FA4E54-C759-480B-A075-CF75E6D85CAD>

[1Z0-070 PDF Dumps](#)

[1Z0-070 VCE Dumps](#)

[1Z0-070 Study Guide](#)