DUMPSARENA

Oracle Exadata X5 Administration

Oracle 1z0-070

Version Demo

Total Demo Questions: 10

Total Premium Questions: 103

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

You are in the process of upgrading your X5 Database Machine half rack to a full rack.

The new storage servers are called DM01CEL08 through DM01CEL14.

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 DATA diskgroup by executing this SQL statement?

A. 0, because the new griddisks will be added to the existing failgroups

B. 12, consisting of seven griddisks each

C. 96, consisting of seven griddisk each

D. 7, consisting of 12 griddisks each

E. 1, consisting of all 96 griddisks

ANSWER: D

Explanation:

The number of failure groups equals the number of Exadata Cells. Each failure group is composed of a subset of grid disks in the Oracle ASM disk group that belong to a single storage cell.

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-administering-asm.htm

QUESTION NO: 2

Which three statements are true about bulk data loading capabilities in an X5 Database Machine environment?

A. DBFS must use the DBFS_DG diskgroup for any DBFS store.

B. DBFS may be used if Exadata-based shared storage is required to stage data before bulk loading into a database.

C. DBFS must be used to have a POSIX-compliant shared storage solution that is accessible from the database servers on a Database Machine.

D. DBFS must be used to bulk load data into a production database on the Database Machine.

E. ACFS may be used if Exadata-based shared storage is required to stage data before bulk loading into a database.

F. ACFS must be used to have a POSIX–compliant shared storage solution that is accessible from the database servers on a Database Machine.

G. ACFS may use the DBFS_DG diskgroup to contain the ADVM volume file.

ANSWER: B C D

Explanation:

External tables on DBFS file-systems provide the probably the most high-performance way to bulk load data into your database.

Bulk Data Loading

Describe the various options for staging data to be bulk loading into Database Machine

Configure the Database File System (DBFS) feature for staging input data files

Use external tables based on input data files stored in DBFS to perform high-performance data loads

QUESTION NO: 3

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 griddisks on 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.

ANSWER: C E

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 NO: 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: B D E

Explanation:

The CellCLI utility supports command-line history.

Each Exadata Storage cell can be monitored, configured, and maintained using the cellcli command line interface.

The dcli utility facilitates centralized management across an Oracle Exadata Storage Server Software realm by automating the execution of CellCLI commands on a set of cells and returning the output to the centralized management location where the dcli utility was run.

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

QUESTION NO: 5

You plan to monitor the status of the motherboard, memory, power, fans, and network cards on the database nodes in your Exadata X6 Database Machine using Enterprise Manager.

Where must you set the thresholds for these hardware components and why, to assure that sensor readings, faults, and any related alerts, are visible in Enterprise Manager?

A. No thresholds need to be set because they are present in the ILOM and in Enterprise Manager.

B. Set thresholds in ILOM and in Enterprise Manager because they are not present anywhere and must be set in both places.

C. Set thresholds only in ILOM because they are not preset anywhere but need to be set only in ILOM.

D. No thresholds need to be set because they are preset in the ILOM and these are sufficient for monitoring.

ANSWER: D

QUESTION NO: 6

Which two Exadata X5 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?

A. Using EXADCLI instead of CELLCLI to create the griddisks

B. Exadata storage realms using database-scoped security mode

C. file permissions on the griddisks in each storage server

D. file permissions on the griddisks in each database server

E. Exadata storage realms using ASM-scoped security mode

ANSWER: B E

Explanation:

Reference: https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmsq/exadatasecurity-features.html#GUID-F53D9493-5927-4106-8D86-65D759419E46

https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmsq/security-guideexadata-database-machine.pdf

QUESTION NO: 7

Which two statements are true about the use of Intelligent Platform Management Interface (IPMI) on an X5 or X6 Database Machine?

A. IPMI can be used for server configuration and management on each database server.

B. The Baseboard Management Controller (BMC), which supports IPMI, runs as a multithreaded process on the storage server O/S.

C. The Baseboard Management Controller (BMC), which supports IPMI, runs as a singlethreaded process on the storage server O/S.

D. IPMI can be used for server configuration and management on each storage server.

ANSWER: A C

Explanation:

A: IPMI – short for Intelligent Platform Management Interface – is an interface standard that allows remote management of a server from another using standardized interface. The servers in the Exadata Database Machine follow that. It's not an Exadata command but rather a general Linux one.

To power on a cell or database server, issue this from another server:

ipmitool -H prolcel01-ilom -U root chassis power on

To stop a server, use the shutdown command. To stop immediately and keep it down, i.e. not reboot, execute:# shutdown -h -y now

C: Like the KCS interface, the SSIF Interface is only specified as a 'Single Threaded Interface' for standard IPMI commands. That is, the BMC implementation is not expected to process more than one IPMI request at a time. While an implementation is allowed to have a degree of 'command queuing', for standard IPMI messages the SSIF lacks a 'Seq' field that software can use to match up particular instances of requests with responses.

References: https://www.intel.com/content/dam/www/public/us/en/documents/specificationupdates/ipmi-intelligent-platform-mgt-interface-spec-2nd-gen-v2-0-spec-update.pdf

QUESTION NO: 8

In which scenario will griddisks residing on the same celldisk have balanced performance characteristics in an X5 High Capacity Exadata Storage Server?

A. griddisks on memory based celldisks

B. griddisks on hard disk-based celldisk consumed by ASM when creating diskgroups using ASM-based Intelligent Data Placement (IDP)

C. griddisks on hard disk-based celldisks consumed by the Database Smart Flash Cache

D. griddisks on Flash-based celldisks

ANSWER: B

Explanation:

Intelligent Data Placement enables you to specify disk regions on Oracle ASM disks for best performance. Using the disk region settings, you can ensure that frequently accessed data is placed on the outermost (hot) tracks which have greater speed and higher bandwidth. In addition, files with similar access patterns are located physically close, reducing latency. Intelligent Data Placement also enables the placement of primary and mirror extents into different hot or cold regions.

References: https://docs.oracle.com/cd/E11882_01/server.112/e18951/ asmdiskgrps.htm#OSTMG10072

QUESTION NO: 9

Which type or types of network traffic are transported over the internal InfiniBand interconnect in Exadata Database Machine X5?

A. IDB protocol traffic, clustered ASM traffic, and clustered database instance traffic

B. IDB protocol traffic and clustered database instance traffic only

C. IDB protocol traffic and clustered ASM instance traffic only

D. Both clustered ASM and clustered database instance traffic only

E. IDB protocol traffic only

ANSWER: E

Explanation:

Oracle Exadata uses the Intelligent Database protocol (iDB) to transfer data between Database Node and Storage Cell Node. iDB is used to ship SQL operations down to the Exadata cells for execution and to return query result sets to the database kernel.

QUESTION NO: 10

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

alter iormplan objective = off alter 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

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.