DUMPSQUEEN

Oracle Exadata X5 Administration

Oracle 1z0-070

Version Demo

Total Demo Questions: 10

Total Premium Questions: 103

Buy Premium PDF

https://dumpsqueen.com

support@dumpsqueen.com

QUESTION NO: 1

Which three are true concerning Exadata snapshot databases?

A. They are supported on non-container databases.

B. They are based on a read-write copy of an existing database.

C. They are integrated with the Multitenant architecture.

D. They can be created only on sparse ASM disk groups.

E. They don't support all Exadata features.

F. They can be created on any type of ASM disk group.

ANSWER: A C D

Explanation:

A: An Exadata snapshot database can be either a non-container database (non-CDB) or a container database (CDB). Creating an Exadata snapshot database of a CDB enables access to all of the pluggable databases in that container.

C: You can create two types of Exadata snapshots, depending on the current setup of your environment:

- You have a pluggable database (PDB) and want to create a test master from it.
- You have a container database (CDB) and want to create test masters from all its PDBs, or you have a simple non-container database and want to create a test master from it.

D: SPARSE disk group based database snapshots is functionality included in Exadata Storage Software. It requires Exadata Storage Software version 12.1.2.1.0 and Oracle Database 12c version 12.1.0.2 with bundle patch 5 or later. This feature is designed to work on native Exadata ASM storage disk groups. It uses ASM SPARSE grid disk based thin provisioning where snapshot databases created on a SPARSE disk group need only the space for changes plus some metadata, thereby enabling storage efficient snapshot databases.

References:

http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-snapshots.htm#SAGUGGUID-E1D6EF45-36EF-40E3-A57E-F80B749E6122

http://www.oracle.com/technetwork/database/exadata/learnmore/exadata-database-copy-twp2543083.pdf

QUESTION NO: 2

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: 3

Identify three valid reasons for creating multiple griddisks on a single harddisk-based celldisk.

A. to implement storage realms so that storage that can be reserved for specific resource consumer groups in the same database

B. to enable the creation of normal or high redundancy ASM diskgroups

C. to segregate storage into multiple pools with different performance characteristics

D. to enable disk mirroring for the system area

E. to implement storage realms so that storage can be reserved for specific databases

F. to implement storage realms so that storage that can be reserved for specific Grid Infrastructure clusters

ANSWER: B C F

Explanation:

Creating multiple grid disks per cell disk allows you to create multiple pools of storage on the same Exadata Storage Server. The multiple grid disks can be assigned to separate ASM diskgroups, which can be provisioned to different databases.

Note: Griddisk is a logical disk that can be created on a celldisk. In a standard Exadata deployment we create griddisks on hard disk based celldisks only. While it is possible to create griddisks on flashdisks, this is not a standard practice.

F: After you complete the cell configuration, you can perform the following optional steps on the storage cell:

- Add the storage cell to the Exadata Cell realm
- Configure security on the Oracle Exadata Storage Server grid disks

References: https://docs.oracle.com/cd/E80920_01/SAGUG/SAGUG.pdf

QUESTION NO: 4

You are patching your Exadata X6 Database Machine by applying a new image to the Storage Servers in a rolling fashion.

Your ASM environment on the Database Machine has five diskgroups stored on an unpartitioned Exadata storage grid, with redundancy settings as shown:

1. DATA_DBM1 – Normal Redundancy

2. RECO_DBM1 – Normal Redundancy

3. DBFS_DG – Normal Redundancy

4. DATA2_DBM1 – High Redundancy

5. DATA3_DBM1 High Redundancy

Which two diskgroups will not suffer from any data loss throughout the patching process even if there is a single disk failure on one of the cells

A. DBFS_DG

B. DATA3_DBM1

C. DATA2_DBM1

D. DATA_DBM1

E. RECO_DBM1

ANSWER: D E

Explanation:

HIGH redundancy provides protection against 2 simultaneous disk failures from 2 distinct storage servers or 2 entire storage servers. HIGH redundancy provides redundancy during Exadata storage server rolling upgrades.

References: http://blog.umairmansoob.com/choosing-high-vs-normal-redundancy-with-exadata/

QUESTION NO: 5

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

alter iormplan objective = low_latency alter iormplan active

There are no database or category plans defined.

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

DSS and Batch workloads perform well some of the time.

Further investigation shows that at times, the workloads are all OLTP I/Os, at other times all batch I/Os, and sometimes a bit of each.

You wish to have disk I/O managed so that performance will be optimized for all workloads.

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

A. alter iormplan objective=high_throughput

B. alter iormplan objective=balanced

C. alter iormplan objective=low_latency

D. alter iormplan objective=auto

E. alter iormplan objective=‘ ’

ANSWER: D

Explanation:

The supported IORM objectives are auto, low_latency, balanced, and high_throughput. 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.

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

QUESTION NO: 6

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: 7

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: 8

Which components of an Exadata storage server image, if updated, may require that a patch also be applied to the database servers in an Exadata X6 Database Machine?

A. Linux operating system on the storage server

B. InfiniBand HCA firmware on the storage server

C. Storage server hard disk device drivers

D. Storage server flash device drivers

ANSWER: D

Explanation:

patches are intended for and include fixes for both the storage servers and Compute servers, and optionally InfiniBand switches.

Patching order

You should patch the Exadata Database Machines in the following sequence

- Oracle GI/RDBMS Homes
- Exadata Storage Cells
- Compute nodes

- Infiniband Switches

References: https://www.toadworld.com/platforms/oracle/w/wiki/11640.oracle-exadata-patching

QUESTION NO: 9

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: 10

Which two are benefits of an active/inactive configured InfiniBand network on Exadata Database Machine X5?

A. Improved performance for Oracle Network traffic

B. Improved reliability for Cache Fusion RAC network traffic

C. Improved reliability when executing Distributed Command Line Interface (DCLI) to run CELLCLI commands

D. Improved performance for ASM rebalance network traffic

E. Improved performance when executing Distributed Command Line Interface (DCLI) to run CELLCLI commands

ANSWER: B C

Explanation:

Active-passive bonding provides reliability through failover.