

Exam Number/Code:1Z0-027

Exam Name:Oracle Exadata
Database Machine X3 Administration

Version: Demo

<http://cert24.com/>

QUESTION NO: 1

To troubleshoot a possible hardware problem, you consider moving all disk drives from one Exadata storage server to a replacement chassis.

You must contain storage availability while performing task.

The Exadata storage server is an X3-8 Database Machine and storage grid is not partitioned.

Which two factors would prevent you from moving the disks from one Exadata storage server to another one?

- A. The existence of an external redundancy ASM diskgroup
- B. The existence of a normal redundancy ASM diskgroup
- C. The existence of an ASM diskgroup with the repair_time attribute set to 0.
- D. The existence of an ASM diskgroup with its compatible.asm attribute set to 10.2.0.0.0
- E. Offline or inactive celldisks in another Exadata server

Answer: A,D

Explanation:

A:If you want Oracle ASM to mirror files, specify the redundancy level as NORMAL REDUNDANCY (2-way mirroring by default for most file types) or HIGH REDUNDANCY (3-way mirroring for all files). You specify EXTERNAL REDUNDANCY if you do not want mirroring by Oracle ASM. For example, you might choose EXTERNAL REDUNDANCY if you want to use storage array protection features.

D:Restoring the redundancy of an Oracle ASM disk group after a transient disk path failure can be time consuming. This is especially true if the recovery process requires rebuilding an entire Oracle ASM failure group. Oracle ASM fast mirror resync significantly reduces the time to resynchronize a failed disk in such situations. When you replace the failed disk, Oracle ASM can quickly resynchronize the Oracle ASM disk extents.

To use this feature, the disk group compatibility attributes must be set to 11.1 or higher.

Incorrect:

Not C:You can set the DISK_REPAIR_TIME disk group attribute to delay the drop operation by specifying a time interval to repair the disk and bring it back online.

Note:

*The redundancy levels are:

/External redundancy

Oracle ASM does not provide mirroring redundancy and relies on the storage system to provide RAID functionality. Any write error cause a forced dismount of the disk group. All disks must be located to successfully mount the disk group.

/Normal redundancy

Oracle ASM provides two-way mirroring by default, which means that all files are mirrored so that there are two copies of every extent. A loss of one Oracle ASM disk is tolerated.

You can

optionally choose three-way or unprotected mirroring.

/High redundancy

Oracle ASM provides triple mirroring by default. A loss of two Oracle ASM disks in different failure groups is tolerated.

Reference:Administering Oracle ASM Disk Groups

QUESTION NO: 2

Which two are regarding the case of storage indexes?

- A. To increase the chance of using the a storage index, you can make table indexes invisible.
- B. To maximize the benefit of storage Indexes, load your data stored on the filtered columns.
- C. The cell physical 10 bytes saved by storage index statistic returns multiple rows, one for each storage server.
- D. Storage indexes are retained after a cell is rebooted.
- E. Avoid the use of bind variables because Storage Indexes do not work with bind variables.

Answer: B,E

Explanation:

B:* To use storage indexes, Oracle Exadata queries must use smart scans, so not all types of applications can benefit from storage indexes.

*With Exadata storage, database operations are handled much more efficiently. Queries that perform table scans can be processed within Exadata storage with only the required subset of data returned to the database server. Row filtering, column filtering and some join processing (among other functions) are performed within the Exadata storage cells. When this takes place only the relevant and required data is returned to the database server.

Incorrect:

Not D:Storage indexes reside in the memory of the storage servers.

If a storage cell is shutdown or rebooted the storage index will be lost from memory and will be recreated on subsequent accesses to the data after the cell has been brought back online.

QUESTION NO: 3

Identify the resource bottleneck for which QoS Management can generate recommendation?

- A. CPU resource bottlenecks
- B. Global Cache resource bottlenecks
- C. I/O resource bottlenecks
- D. Network resource bottlenecks

Answer: B

Explanation:

Note:

*Oracle Exadata QoS Management provides the following benefits:

?Reduces the time and expertise requirements for system administrators who manage Oracle

Real Application Clusters (Oracle RAC) resources

?Helps reduce the number of performance outages

?Reduces the time needed to resolve problems that limit or decrease the performance of your applications

?Provides stability to the system as the workloads change

?Makes the addition or removal of servers transparent to applications

?Reduces the impact on the system caused by server failures

?Helps ensure that service-level agreements (SLAs) are met

?Enables more effective sharing of hardware resources

?Protects existing workloads from over committed memory-induced server failures

?Exadata Storage Virtualization

?Exadata provides a rich set of sophisticated and powerful storage management virtualization capabilities that leverage the strengths of the Oracle Database, the Exadata software, and Exadata hardware.

*When QoS Management is enabled, individual Exadata nodes are protected from memory related failures. The Memory Guard feature tracks real-time memory use and should it detect a node has over-committed memory, will prevent new database requests from being sent until the current load is relieved

QUESTION NO: 4

Which two are Oracle recommendations for media based backups performed for a database running on a Database Machine?

A. Allocate equivalent number of channels and instances per tape drive.

B. Perform periodic level 0 backups and daily cumulative level-1 backups.

C. Use the InfiniBand network between the database server and media server.

D. Configure Recovery Manager (RMAN) channels to connect to the least loaded instances.

E. Use InfiniBand network between the media server and the storage servers.

Answer: A,C

Explanation:

A:Configure one RMAN channel per tape drive and add tape drives to scale backup rates.

C:Configure the Preferred Network Interface (PNI) to direct the Oracle Secure Backup traffic over the InfiniBand network interface.

Example:

ob> lspni (List Preferred Network Interface)

mediaserver1:

PNI 1:

interface: mediaserver1-ib

clients: dbnode1, dbnode2, dbnode3, dbnode4,
dbnode5, dbnode6, dbnode7, dbnode8

PNI 2:

interface: mediaserver1

clients: adminserver

dbnode1:

PNI 1:

interface: dbnode1-ib

clients: mediaserver1

Note:

*Using the Sun ZFS Backup Appliance as an Oracle RMAN backup target for an Oracle Exadata system delivers much faster backup and recovery, enabling organizations to achieve shorter recovery time objectives and shrink backup windows. The appliance is designed for high sustained read and write I/O performance, and it is connected to the Oracle Exadata system via a highthroughput InfiniBand network fabric.

*As the only unified storage vendor to support InfiniBand as a storage network for backup and restore operations, Oracle is leading the way with native high-bandwidth interconnects.

*The InfiniBand network provides 40 Gb of bandwidth per port between the database servers, storage cells, and the Sun ZFS Backup Appliance. Backup and restore operations can be automatically parallelized across all database nodes, Oracle Exadata storage cells, Sun ZFS Backup Appliance channels, and controllers.

QUESTION NO: 5

What is the benefit of bonding the client access network configuration?

- A. Improved performance
- B. Improved reliability
- C. Both improved performance and reliability
- D. A Single Client Access Name (SCAN)
- E. Improved monitoring

Answer: C

Explanation:

The network ports can be connected directly between the components or the ports can be bonded together to form a logical interface. The bonded ports can be configured for load balancing, fault tolerance, and better network utilization

Reference:Oracle Exadata Database Machine -Backup & Recovery Sizing:Tape Backups

QUESTION NO: 6

To guarantee proper cooling, you plan to place perforated floor tiles near your Database Machine.

Where, in relation to the cabinet, should they be placed?

- A. On the left side, because the air flow is from left to right
- B. At the back, because the air flow is from back to front
- C. On the right side, because the air flow is from right to left.
- D. At the front, because the air flow is from front to back
- E. Underneath the cabinet, because the air flow is from bottom to top

Answer: D

Explanation: Airflow must be front-to-back.

Reference: Oracle White Paper,ORACLE Exadata DatabaseMachine X3-8

QUESTION NO: 7

Identify two valid reasons for creating multiple griddisks on a Single celldisk.

- A. To segregate storage into multiple pools with different performance characteristics
- B. To facilitate normal or high redundancy ASM diskgroups
- C. To enable disk mirroring for the system area
- D. To segregate storage into multiple pools that can be assigned to different databases
- E. To segregate storage into multiple pools that can be assigned to different resource consumer groups in the same database.

Answer: B,D

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:

*Celldisks are the third layer of abstraction. It was introduced to enable interleaving in the first place

*Griddisks are the fourth layer of abstraction, and they will be the Candidate Disks to build your ASM diskgroups from.

* The first grid disk created on the cell disk will allocate space from the outer tracks and move towards the inner tracks, reserving the number of tracks that correspond to the size of the grid disk.This grid disk provides the fastest performance since the outer tracks of a

hard disk provide the best read/write performance. The next grid disk you will create starts from the tracks where the first grid disk ends, and this process repeats until you exhaust all the space on the cell disk or you are done creating the grid disks.

QUESTION NO: 8

You plan to migrate an Oracle database that supports an online transaction processing (OLTP) workload to your Database Machine.

Following are details for the source database:

Database version: 10.2.0

Byte order: Big Endian HP-UX (64-bit)

Database size: 24 TB

Storage: ASM with 1 MB allocation unit size

Which two are supported migration methods?

- A. Physical migration using ASM online Migration
- B. Physical migration using Transportable Database
- C. Logical migration using Oracle Streams
- D. Local migration using Oracle Streams
- E. Logical migration using logical standby

Answer: B,C

Explanation:

Note:

* physical migration techniques (e.g. Data Guard Physical Standby, Transportable Database, Transportable Tablespaces(B), Rapid Clone) vs. logical migration techniques(C)(e.g. Oracle Data Pump).

*You can use either Data Pump export/import or original export/import for Streams instantiations.

*The Oracle Exadata Database Machine runs the standard Oracle Database. Therefore, any application that runs with the Oracle Database today can be seamlessly migrated to the Exadata Database Machine with no changes to the application.

*There are several techniques for migrating data to a Database Machine. Migration can be done using Oracle Recovery Manager (RMAN) to backup from traditional storage and restore the data onto Exadata. Oracle Data Guard can also be used to facilitate a migration. This is done by first creating a standby database based on Exadata storage. The standby can be using Exadata storage and the production database can be on traditional storage. By executing a fast switchover, taking just seconds, you can transform the standby database into the production database. This provides a built-in safety net as you can undo the migration very gracefully if unforeseen issues arise.

Transportable Tablespaces(B) and Data Pump may also be used to migrate to Exadata. Any technique used to move data between Oracle Databases can be used with Exadata.

QUESTION NO: 9

You have configured a multi-rack Database Machine with two X3-8 full racks all in a single cluster and storage grid.

Which two are true regarding the servers on which Enterprise manager agents must be deployed in order to monitor all components of the multi-rack Database Machine?

- A. On only one database server in the first rack
- B. On all storage servers in all racks
- C. On at least two storage servers in the first rack
- D. On all database servers in the first rack
- E. On only one database server in second rack
- F. On all database servers in second rack
- G. On at least two storage servers in the second rack

Answer: D,F

Explanation:

Note:

*The Enterprise Manager agent must be deployed to all compute nodes of the Database Machine.

*Oracle's documentation uses the term compute node when referring to the database server tier of the platform.

*The Exadata Database Machine runs Oracle Database 11g Real Application Cluster. The cluster and the database run on the servers known as database nodes or compute nodes (or simply "nodes").

*Cells and compute nodes are not shared between partitions.

*Compute nodes in same partition share the same Cluster.

QUESTION NO: 10

Which three storage components are available after the standard initial Database machine deployment?

- A. The DATA_<DBM_Name> ASM diskgroup
- B. The RECO_<DBM_Name> ASM diskgroup
- C. Mirrored system partitions on hard disk 0 and hard disk 1
- D. The DBFS_DG diskgroup with external redundancy
- E. Exadata Smart Flash Cache using all of the flashdisk space

Answer: A,B,D

Explanation:

The default configuration on Exadata is to have 3 diskgroups - DATA, RECO, and

DBFS_DG.