

Oracle

Exam 1z0-027

Oracle Exadata Database Machine X3 Administration

Version: 6.13

[Total Questions: 72]

Question No : 1

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 high-throughput 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 : 2

Last weekend, an Exadata storage server flashdisk entered the predictive failure state.

The flashdisk is used by the flashcache and has a griddisk which is a member of a normal redundancy diskgroup.

Identify the four steps you must perform to replace this flashdisk.

- A. Identify the griddisk on the predictive failure flashdisk and drop it from the associated ASM diskgroup
- B. Verify that the griddisk located on the predictive failure flashdisk has been successfully dropped from the associated ASM diskgroup.
- C. Drop the flashcache on the cell and re-create it using all but the predictive failure flashdisk.
- D. Safely power off the cell containing the predictive failure flashdisk.
- E. Replace the predictive failure flashdisk.
- F. Power up the cell containing the replaced flashdisk and activate all griddisks.
- G. Drop the flashcache on the cell and re-create it using all flashdisks.
- H. Create a new griddisk on the replaced flashdisk.
- I. Add the griddisk back into the ASM diskgroup to which it belonged.

Answer: A,D,E,F

Question No : 3

Your Database Machine has a large database with some very large tables supporting OLTP workloads.

High volume Insert applications and high volume update workloads access the same tables.

You decide to compress these tables without causing unacceptable performance overheads to the OLTP application.

Which three are true regarding this requirement?

- A. Using 'compress for oltp' will compress the data less than if using Hybrid Columnar Compression when specified with compress for query low.
- B. The compression is performed on the storage servers when using compress for oltp in an Exadata environment.
- C. The compression method compress for archive high is the worst fit for this requirement.
- D. Using 'compress for oltp' will compress the data more than if using Hybrid Columnar Compression when specified with compress for archive low.
- E. The compression is performed on the database servers when using 'compress for oltp' in an Exadata environment.

Answer: A,C,E

Explanation:

Note:

(E not B):

* Types of compression

Basic compression

OLTP compression

Warehouse compression

Online archival compressio

*

/ OLTP compression allows compression during DML operations.

/ Basic compression works at the data block level.

* When you enable table compression by specifying COMPRESS FOR OLTP, you enable OLTP table compression. Oracle Database compresses data during all DML operations on the table. This form of compression is recommended for OLTP environments.

* When you specify COMPRESS FOR QUERY or COMPRESS FOR ARCHIVE, you enable hybrid columnar compression. With hybrid columnar compression, data can be compressed during bulk load operations. During the load process, data is transformed into a column-oriented format and then compressed. Oracle Database uses a compression algorithm appropriate for the level you specify. In general, the higher the level, the greater

the compression ratio.

Hybrid columnar compression can result in higher compression ratios, at a greater CPU cost. Therefore, this form of compression is recommended for data that is not frequently updated.

Question No : 4

Which two are true about Exadata storage server alerts?

- A. Metric alerts are never stateful.
- B. Metrics have no thresholds set on them by default.
- C. SNMP alert notifications can be sent to only one destination.
- D. Metric threshold; if set, will persist across storage server reboots.
- E. SMTP alert notifications must be sent to both ASR manager and Enterprise Manager Agents

Answer: D,E

Explanation:

Incorrect:

Not A: there are both stateful and stateless alerts.

Not B: Metrics have defaults.

Not C: An SNMP alert can have multiple destinations.

Question No : 5

Which three statements are true regarding the configuration of Auto Service Request (ASR) on your Database Machine?

- A. HTTPS connectivity must be enabled from the ASK Manager host to the internet.
- B. ASR Manager must be installed on one of the database servers.
- C. Oracle Solaris must be used on the server where ASK Manager is installed.
- D. SMTP must be enabled on at least one database server.
- E. Oracle Linux must be used on the server where ASR Manager is installed

- F. ASR Manager may be installed on any type of server running Oracle Linux.
 G. ASR Manager may be installed on any type of server running Oracle Solaris.

Answer: A,F,G

Explanation: A: Before installing ASR, please ensure the following conditions are met:

/ (A) Ensure connectivity to the Internet using HTTPS.

/ Make sure you have access to My Oracle Support and that your contact information is correct and current.

/ Make sure all of your assets have a Contact assigned and that the contact is correct and current.

/ Identify and designate a system to serve as ASR Manager.

/ Identify and verify ASR assets.

F, G (not C, not E, not B): The recommended configuration is to install the ASR Manager, which receives fault telemetry information from the servers in Oracle Exadata Database Machine, on an external standalone server. This server must run Solaris or Linux as the operating system.

Reference: Oracle Auto Service Request Exadata Database Machine Quick Installation Guide

Question No : 6

You are checking the status of the ports on the one of the InfiniBand switches in your Database Machine.

You have run the getportstatus as shown:

```
[root@exampsw-ib2 bin]# getportstatus 17b
Port status for connector 17B Switch Port 2
Adminstate:.....Enabled
LinkWidthEnabled:.....1X or 4X
LinkWidthSupported:.....1X or 4X
LinkWidthActive:.....4X
LinkSpeedSupported:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkState:.....Active
PhysLinkState:.....LinkUp
LinkSpeedActive:.....10.0 Gbps
LinkSpeedEnabled:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
```

You have 36 ports to check.

Which two methods would you use to check the status of all 36 ports?

- A. Use Enterprise Manager.
- B. Use the `ibqueryerros.pl` script logged into the InfiniBand switch as root.
- C. Use Database Control.
- D. Create user-defined metrics for the InfiniBand switch.

Answer: A,B

Explanation: http://docs.oracle.com/cd/E17559_01/install.111/e20085/e20085-03.htm

Question No : 7

Which four statements are true about Exadata Smart Flash Cache?

- A. Smart Scan will always be done for I/Os to flash based griddisks.
- B. Flash based ASM diskgroups may share space with the Flash Cache on the flashdisks.
- C. Single block reads can benefit from Smart Flash Cache.
- D. Smart Scan will never be done for I/Os to flash based griddisks.
- E. Multiblock reads can benefit from Smart Flash Cache.
- F. Smart Flash Logs reduce the size of Smart Flash Cache.

Answer: B,C,E,F

Explanation: B: * Grid disks (the logical disks

that reside on physical cell disks) are created on these flash-based cell disks and the grid disks are assigned to an Automatic Storage Management (ASM) diskgroup. The best practice would be to reserve the same amount of flash on each Exadata cell for flash disks and have the ASM

diskgroup spread evenly across the Exadata cells in the configuration just as you would do for

regular Exadata grid disks. This will evenly distribute the flash I/O load across the Exadata cells

and flash.

Note:

* The Exadata

Storage Server comes with a substantial amount of flash storage. A small amount is allocated for

database logging and the remainder will be used for caching user data

Question No : 8

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

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

- A. Set thresholds in ILOM and in Enterprise Manager because they are not preset anywhere and must be set in both places
- B. Set thresholds only in ILOM because they are not preset anywhere but need only be set in ILOM.
- C. No thresholds need be set because they are preset in the ILOM and in Enterprise Manager.
- D. No thresholds need ho set because they are preset in the ILOM and these are sufficient for monitoring.

Answer: D

Explanation:

Note:

* The ILOM management interface is also integrated with Oracle Enterprise Manager Ops Center.

* ILOM Features and Functionality include:

Remote hardware monitoring which include:

Monitor customer-replaceable units (CRUs) and fieldreplaceable units (FRUs), including power supplies, fans, host bus adapters (HBAs), disks, CPUs, memory, and motherboard

* The ILOM-based service processor (SP) receives error telemetry about error events that occur within the major system components on the host (CPU, memory, and I/O hub) and the environmental subsystem within the chassis (such as fans, power supplies, and temperature). The components and conditions are then diagnosed as fault events and captured in the ILOM event log.

Question No : 9

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: A,B

Explanation: 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.

Question No : 10

Which three factors should you consider when choosing a method for migrating a database to the Database Machine?

- A. Type of database workload
- B. Number of tablespaces in the source database
- C. Size of the source database
- D. Endian format of the source database
- E. ASM attributes and allocation unit (AU) size of the diskgroups used by the source database

Answer: A,C,D

Explanation: A: Use Real Production Workload

– Real Application Testing (RAT)

- Database Workload Replay

- SQL Performance Analyzer (SPA)

D:

- * Endian format: Non-Exadata format

Reference: Exadata MAA Best Practices, Migrating Oracle Databases

Question No : 11

Which three are true about Smart Flash log?

- A.** I/O Resource Manager database plans can be used to enable or disable Smart Flash Log for different databases.
- B.** LGWR will not wait for writes to Smart Flash log if the write to a disk based log file completes first.
- C.** Smart Flash Log is enabled by default, using 1024 MB of Flash storage on each storage server.
- D.** You can remove Smart Flash Log from a single storage server with the drop flashing command.
- E.** I/O Manager category plans can be used to enable or disable Smart Flash Log For different I/O categories.

Answer: A,B,D

Explanation: A: the Exadata I/O Resource Manager (IORM) has been enhanced to enable or disable Smart Flash Logging for the different databases running on the Database Machine.

B: Smart Flash Logging works as follows. When receiving a redo log write request, Exadata will do parallel writes to the on-disk redo logs as well as a small amount of space reserved in the flash hardware. When either of these writes has successfully completed the database will be immediately notified of completion. If the disk drives hosting the logs experience slow response times, then the Exadata Smart Flash Cache will provide a faster log write response time. Conversely, if the Exadata Smart Flash Cache is temporarily experiencing slow response times (e.g., due to wear leveling algorithms), then the disk drive will provide a faster response