1Z0-027 Q&As

Oracle Exadata X3 and X4 Administration

Pass Oracle 1Z0-027 Exam with 100% Guarantee

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

https://www.pass4lead.com/1Z0-027.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle Official Exam Center

QUESTION 1

You are using Hybrid Columnar Compression for a table stored in a tablespace that is contained in an Exadata-based ASM diskgroup. Identify three statements that correctly explain where the compression and decompression can be done.

A. Decompression can be done on the database servers.

B. Compression can be done on the Exadata storage servers.

C. Compression can be done on the database servers.

D. Decompression can be done on the Exadata storage servers.

Correct Answer: ABD

Explanation: B:

*Exadata storage provides an advanced compression technology, Hybrid Columnar Compression, that typically provides 10x, and higher, levels of data compression.

*The Exadata Storage Server (Exadata storage or Exadata cells) is used as the storage for the Oracle Database in the Database Machine. It runs the Exadata Storage Server

Software that provides the unique and powerful Exadata technology including Smart Scan, Smart

Flash Cache, Smart Flash Logging, IO Resource Manager, Storage Indexes and Hybrid Columnar Compression.

A, D:

*decompression

/Queries run directly on Hybrid Columnar Compressed data does not require the data to be decompressed

/Data that is required to satisfy a query predicate does not need to be decompressed; only the columns and rows being returned to the client are decompressed in memory

/The decompression process typically takes place on the Oracle Exadata Storage Server in order to maximize performance and offload processing from the database server.

QUESTION 2

Which type of network traffic is transported over the internal InfiniBand network in a Database Machine?

A. IDB protocol traffic only

B. Both Clustered ASM and RAC database instance traffic

C. Clustered ASM Instance traffic only

D. RAC database instance traffic only

E. IDB protocol traffic,Clustered ASM traffic, and RAC database instance traffic

Correct Answer: A

Explanation: The Exadata software is optimally divided between the database servers and Exadata

cells. The database servers and Exadata Storage Server Software communicate using the iDB ?

the Intelligent Database protocol. iDB is implemented in the database kernel and transparently

maps database operations to Exadata-enhanced operations. iDB implements a function shipping

architecture in addition to the traditional data block shipping provided by the database. iDB is

used to ship SQL operations down to the Exadata cells for execution and to return query result

sets to the database kernel. Instead of returning database blocks, Exadata cells return only the

rows and columns that satisfy the SQL query. Like existing I/O protocols, iDB can also directly

read and write ranges of bytes to and from disk so when offload processing is not possible Exadata operates like a traditional storage device for the Oracle Database. But when feasible, the

intelligence in the database kernel enables, for example, table scans to be passed down to execute

on the Exadata Storage Server so only requested data is returned to the database server. iDB is built on the industry standard Reliable Datagram Sockets (RDSv3) protocol and runs over

InfiniBand. ZDP (Zero-loss Zero-copy Datagram Protocol), a zero-copy implementation of RDS, is used to eliminate unnecessary copying of blocks. Multiple network interfaces can be used on the database servers and Exadata cells. This is

an extremely fast low-latency protocol that minimizes the number of data copies required to service I/O operations.

Note:

*The Database Machine uses a state of the art InfiniBand interconnect between the servers and

storage. Each database server and Exadata cell has dual port Quad Data Rate (QDR) InfiniBand

connectivity for high availability.

*The same InfiniBand network also provides a high performance cluster interconnect for the Oracle Database Real Application Cluster (RAC) nodes.

Note:

*An InfiniBand network allows you to connect multiple Oracle Exadata Database Machines to form a larger single system image configuration; each InfiniBand link provides 40 Gigabits of bandwidth?any times higher than traditional storage or server networks

QUESTION 3

You plan to migrate your Oracle Version 11.1.0.2 database to your Exadata Database Machine.

The database supports an online transaction processing (OLTP) workload and is currently hosted on a Little Endian platform

Which two are the supported and appropriate migration methods to minimize downtime?

A. Upgrade source database to 11.2.0 and migrate using a physical standby database.

B. Migrate using Data Pump.

C. Migrate using GoldenGate.

D. Migrate using cross platform Transportable Database.

E. Migrate using ASM online migration.

Correct Answer: BD

Explanation: B:Oracle Data Pump can be used to migrate an Oracle database to a new platform, and to move from an older release of the database to a newer release. Using Oracle Data Pump to move an E-Business Suite database is a well-documented and tested procedure, and can be used to migrate your database to the Oracle Exadata Database Machine and upgrade it to Oracle Database 11g release 2 (11.2) in the same exercise

Reference:Migrating the Oracle E-Business Suite Database to Oracle Exadata Database Machine UsingTransportable Tablespaces

Reference:Migrating Oracle E-Business Suite to Oracle Exadata Database Machine Using Oracle Data Pump

QUESTION 4

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.

Correct Answer: ACE

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 5

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 thefoursteps 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 flashdlsk.

E. Replace the predictive failure flashdisk.

F. Power up the cell containing the replaced flashdlsk and activate all grlddlsks.

G. Drop the flashcache on the cell and re-create it using all flashdlsks.

H. Create a new griddisk on the replaced flashdisk.

I. Add the griddisk back into the ASM diskgroup to which it belonged.

Correct Answer: ADEI

Note:

*Exadata monitors for the number of media and other disk/flash failures (e.g. an I/O write failure due to physical media damage). If there are too many of those, Exadata is \\'predicting\\' that it will soon fail and it takes it out of the system.

*Exadata Server, that runs on the storage cells, monitors disk health and performance. If the disk performance degrades it can put it into proactive failure mode. It also monitors for predictive failures based on the disk\\'s SMART (Selfmonitoring, Analysis and Reporting Technology) data. In both cases, the Exadata Server notifies XDMG to take those disks offline.

When a faulty disk is replacedf on the storage cell, the Exadata Server will recrate all grid disks on a new disk. It will then notify XDMG to bring those grid disks online or add them back to disk groups, in case they were already dropped.

*ASM is a critical component of the Exadata software stack. It is also a bit different - compared to non-Exadata environments. It still manages your disk groups, but builds those with grid disks. It still takes care of disk errors, but also handles

predictive disk failures. It doesn\\'t like external redundancy and ACFS, but it makes the disk group smart scan capable.