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.pass4itsure.com/1Z0-027.html

100% Passing Guarantee
100% Money Back Assurance

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

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

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

Correct Answer: DE

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 2

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 3

You have used setupem.sh to deploy a plug in for Grid Control.

In which two ways do all the targets supported by the plug-in get configured?

A. The targets must be configured by an Enterprise Manager administrator manually using Grid Control.

B. Setupem.sh is used again to configure the targets.

C. The targets must be configured by an O/S administrator manually.

D. The targets must be configured by a database administrator manually.

E. There may be more than one target for each plug-in.

Correct Answer: AE

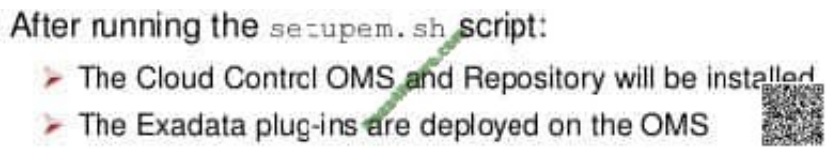

QUESTION 4

You plan to monitor storage servers after configuring an I/O resource manager plan with directives for inter-database plans and intra-database plans. Which two types if metrics would help assess the impact of the intra-database plans on I/O to the storage servers?

A. Category I/O

B. Database I/O

C. Resource Consumer Group I/O

D. Smart Flash Log I/O

E. Smart Flash Cache I/O

Correct Answer: BC

Explanation: B:Database metrics provide information about the size of the I/O load from each database specified in the interdatabase plan.

C:Consumer group metrics provide information about the size of the I/O load from each consumer group specified in a database resource plan. Each database in the interdatabase plan has metrics for each of its consumer groups.

Note:

*I/O Resource Manager (IROM) Settings

Incorrect:

Not A:Category metrics provide information about the size of the I/O load from each category specified in the current IORM category plan.

QUESTION 5

You are evaluating the performanceof a SQLstatementthat accessesa very large table, and have run the following queryproducingthe output shown:

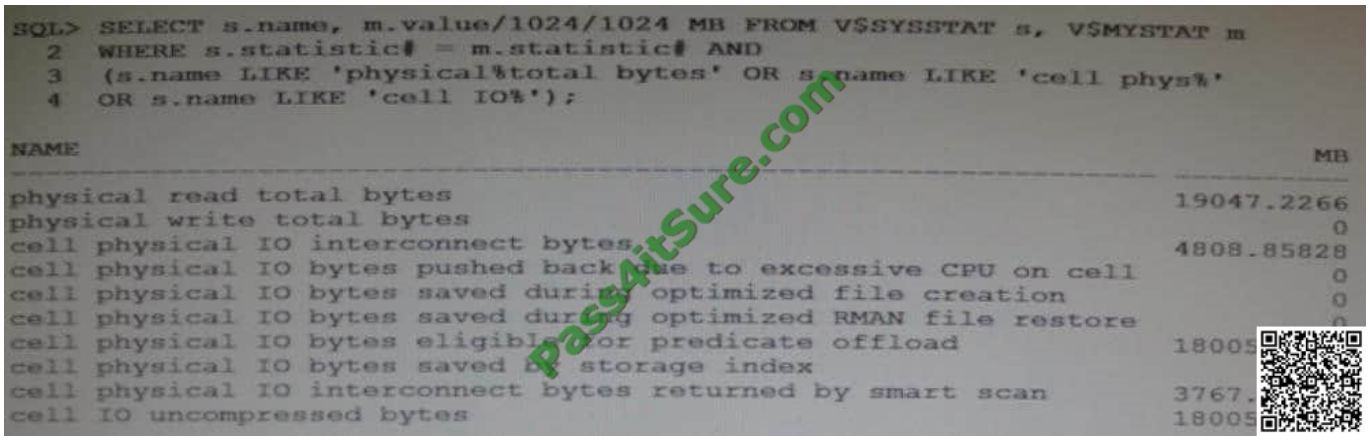

For which two reasons would the; "physical read total bytes" statistic be greater than the "cell physical IO bytes eligible for predicate offload" statistic?

A. There is an index on the column used in the where clause, causing "cell multiblock physical reads" to be requested

by the database instance, resulting in additional I/O for blocks in the cells.

B. The table is an IOT and has an overflow segment, causing "cell multiblock physical reads" to be requested by the database instance, resulting in additional I/O for block in the cells.

C. There is an uncommitted transaction that has modified some of the table blocks, causing some "cell single block physical reads" to be requested by the database instance, resulting in additional I/O for block in the cells.

D. The table is an indexed clustered table, causing "cell single block physical reads" to be requested by the database instance, resulting in additional I/O forblocks in the cells.

E. There are migrated rows in the table, causing some "cell single block physical reads" to be requested by the database instance, resulting in additional I/O for blocks in the cells.

Correct Answer: BE

Note:

* physical read total bytes: the size of the segment to read is known by the database, and must be read entirely from the database\\'s perspective. *cell physical IO bytes eligible for predicate offload: this statistic shows the amount of data which the cell server is able to process on behalf of the database, instead of the database processing and the cell server just delivering blocks. *Cell physical IO bytes eligible for predicate offload --- This number should be high The higher the number more MB/GB is filtered out at the cell level itself rather sending it to the buffer cache to filter the rows.

*In this case, all bytes are processed on the cellserver (cell physical IO bytes eligible for predicate offload=physical read total bytes)

*Cell Offloading:The storage cells are intelligent enough to process some workload inside them, saving the database nodes from that work. This process is referred to as cell offloading.

QUESTION 6

A table in one of your database schemas contains only varchar, number, and date data types for the columns.

Which three operations can be offloaded to the Exadata storage servers when doing a smart scan against this table, if no other situations arise that prevent Smart Scan from occurring?

A. Column filtering

B. Sort merge join filtering

C. Predicate filtering

D. Nested loop Join filtering

E. Hash join filtering

F. Virtual column filtering

Correct Answer: ACF

Explanation: Exadata 11.1 Smart Scan operations include:

- Restriction (filtering of rows)

- Projection (filtering of columns)

- Join Filters (Bloom Filters) commonly used in Fact/Dimension joins found in Star Schemas. This is different than Star Transformation. Other join operations such as HASH(not E), SORT, SORT-MERGE(not B), NESTED LOOP(not D), etc. are done by the Oracle Database Grid.

A:Smart Scan Column Filtering Exadata provides column filtering, also called column projection, for table scans. Only the columns requested are returned to the database server rather than all columns in a table. For example, when the following SQL is issued, only the employee_name and employee_number columns are returned from Exadata to the database kernel. SELECT employee_name, employee_number FROM employee_table. For tables with many columns, or columns containing LOBs (Large Objects), the I/O bandwidth saved can be very large. Using both predicate and column filtering dramatically improves performance and reduces I/O bandwidth consumption. In addition, column filtering also applies to indexes, allowing for even faster query performance.

C:Smart Scan Predicate Filtering

Exadata enables predicate filtering for table scans. Only the rows requested are returned to the database server rather than all rows in a table. For example, when the following SQL is issued only rows where the employees\\' hire date is after the specified date are sent from Exadata to the database instance.

SELECT * FROM employee_table WHERE hire_date > \\'1-Jan-2003\\'. This ability to return only relevant rows to the server greatly improves database performance. This performance enhancement also applies as queries become more complicated, so the same benefits also apply to complex queries, including those with subqueries.

QUESTION 7

You are about to replace one memory DIMM in an Exadata storage server and need to power off the affected cell.

Which four commands must you execute to safely power off the storage server in your standard deployed quarter rack Database Machine assuming that redundancy is not compromised?

A. `crsctl stop cluster -all\\' on one of the database servers

B. CellCLI> LIST GRIDDISK ATTRIBUTE name WHERE asmdeactivationoutcome != `Yes\\'

C. CellCLI> ALTER GRIDDISK ALL INACTIVE

D. CellCLI> LIST GRIDDISK WHERE STATUS != `inactive\\'

E. `shutdown ? now\\' on the affected cell

Correct Answer: BCDE

B: Step 1:

Run the following command to check if there are other offline disks CellCLI> LIST GRIDDISK ATTRIBUTES name WHERE asmdeactivationoutcome != \\'Yes\\'

If any grid disks are returned, then it is not safe to take the storage server offline because proper Oracle ASM disk group redundancy will not be intact.

C: Step 2:

Inactivate all the grid disks when Oracle Exadata Storage Server is safe to take offline using the following command:

CellCLI> ALTER GRIDDISK ALL INACTIVE

Taking the storage server offline when one or more grid disks are in this state will cause Oracle ASM to dismount the affected disk group, causing the databases to shut down abruptly.

D: Step 3: -- Verify all grid disks are INACTIVE to allow safe storage server shut down by running the following command:CellCLI> LIST GRIDDISK ATTRIBUTES name, asmmodestatusCellCLI> LIST GRIDDISKIf all grid disks are INACTIVE, then the storage server can be shutdown without affecting database availability

E:Step 4:

To stop a server, use the shutdown command. To stop immediately and keep it down, i.e.

not reboot, execute:

#

shutdown -h -y now

QUESTION 8

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

Correct Answer: AD

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 9

Which tool will provide you with diagnostic information for all the software log, trace files, and OS information on Database Machine?

A. dbmcheck.sh

B. diagget.sh

C. oswatcher

D. adrci

E. Enterprise Manager

Correct Answer: B

Explanation: Gather all diagnostics information /opt/oracle.SupportTools/onecommand/diagget.sh

QUESTION 10

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.

Correct Answer: BD

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 11

Identify the three components that serve a purpose only in the Database Machine.

A. ASM intelligent Data Placement (IDP)

B. Intelligent Database Protocol (IDB)

C. Database Resource Manager (DBRM)

D. I/O Resource Manager (IORM)

E. Database Filesystem (DBFS)

F. The DISKMON process

Correct Answer: ABD

Explanation: A:Intelligent Data Placement, a feature of ASM that allows placing data in such a way that more frequently accessed data is located close to the periphery of the disk where the access is faster.

B: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

D:The inter-database I/O allocations are defined within the software in the Exadata cell and managed by the I/O Resource Manager (IORM). The Exadata cell software ensures that inter-database I/O resources are managed and properly allocated within, and between, databases.

QUESTION 12

Which two are regarding the case of storageindexes?

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.

Correct Answer: BE

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.

To Read the Whole Q&As, please purchase the Complete Version from Our website.

Try our product !

100% Guaranteed Success
100% Money Back Guarantee
365 Days Free Update
Instant Download After Purchase
24x7 Customer Support
Average 99.9% Success Rate
More than 800,000 Satisfied Customers Worldwide
Multi-Platform capabilities - Windows, Mac, Android, iPhone, iPod, iPad, Kindle

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications.
You can view Vendor list of All Certification Exams offered:

https://www.pass4itsure.com/allproducts

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket: