

Exam Questions 1Z0-027

Oracle Exadata Database Machine Administration, Software Release 11.x

<https://www.2passeasy.com/dumps/1Z0-027/>



NEW QUESTION 1

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.

Answer: AE

NEW QUESTION 2

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 Database Machine X3-8

NEW QUESTION 3

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

Answer: B

Explanation:

Gather all diagnostics information

/opt/oracle.SupportTools/onecommand/diagget.sh

NEW QUESTION 4

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: AD

NEW QUESTION 5

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

Answer: E

NEW QUESTION 6

Identity 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: A

NEW QUESTION 7

You have two very large databases supporting OLTP workloads which run on multiple small-blade style servers in a cluster and which require optimal latency for I/O.

You plan to migrate them to a Database Machine once a capacity planning exercise is finished.

The database backup strategy requires that backups are written directly to media.

High availability requirements state that you must be able to survive node failures at any time.

Which three Database Machine components or features would you recommend to support these requirements?

- A. Use of write back flash cache
- B. Use of smart flash logs
- C. High capacity disks in the Database Machine
- D. High performance disks in the Database Machine
- E. A high capacity expansion full rack
- F. A high performance expansion full rack

Answer: ABC

NEW QUESTION 8

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: DE
Explanation: Incorrect:

Explanation:

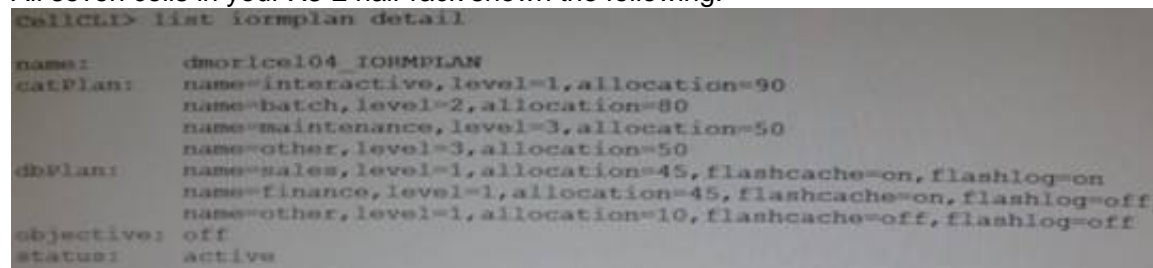
Not A: there are both stateful and stateless alerts. Not B: Metrics have defaults.

Not C: An SNMP alert can have multiple destinations.

NEW QUESTION 9

You are examining the existing IORM configuration on the cells of Database Machine, to see if they require my modifications based on recent changes to various workloads.

All seven cells in your X3-2 half-rack shown the following:



```
CellCLI> list iormplan detail
name:          dmoricel04_IORMPLAN
catPlan:       name=interactive,level=1,allocation=90
               name=batch,level=2,allocation=80
               name=maintenance,level=3,allocation=50
               name=other,level=3,allocation=50
dbPlan:        name=sales,level=1,allocation=45,flashcache=on,flashlog=on
               name=finance,level=1,allocation=45,flashcache=on,flashlog=off
               name=other,level=1,allocation=10,flashcache=off,flashlog=off
objective:     off
status:        active
```

Which two are true about I/O to the cells using this plan?

- A. I/O requests in the batch category may use flashcache if the I/O is from the sales finance database, and these I/O requests are guaranteed to get 80% of the I/O if the interactive category I/Os use no more than 20%.
- B. I/O requests made by sessions in the marketing database may use flashing and flashcache if no other categories or database or database are using flashing and flashcache at the same time.
- C. I/O requested in the interactive category may use flashdns if the I/O is from the sales or finance databases, and these I/O requests are guaranteed to get 90% of the I/O if the enough I/Os are issued in this category.
- D. I/O requests from the sales database may use flashing regardless of the I/O category.
- E. No I/Os in any category or from any database may use flashing or flashcache because the objective is off.

Answer: AD

NEW QUESTION 10

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

Answer: ACE

Explanation:

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.

NEW QUESTION 10

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: BCEF

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

NEW QUESTION 15

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: DF

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.

NEW QUESTION 18

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

Answer: ABD

Explanation:

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.

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

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.

NEW QUESTION 23

Your database Machine has the exachk utility pre-installed and you decide to use it periodically, to validate the installation against Oracle's recommended best practices.

Which two actions could you take to do this?

- A. Use a cron job on a database node to run it at regular intervals.
- B. Run it once from a database node and it will then perform periodic monitoring automatically.
- C. Use a cron job on each cell to run it at regular intervals.
- D. Run it once on each cell and it will then perform periodic monitoring automatically.
- E. Create a Job in Enterprise Manager to run the exachk utility at regular intervals.

Answer: AE

NEW QUESTION 24

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: CE

Explanation:

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.

NEW QUESTION 29

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: 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

NEW QUESTION 32

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: B

Explanation:

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

NEW QUESTION 33

Which three are true about Enterprise Manager plug-in configuration for the Database Machine?

- A. There are several separate plug-ins for Grid Control.
- B. There is one plug-in for Cloud Control.
- C. There are several separate plug-ins for Cloud Control.
- D. Some plug-ins require SNMP trap forwarders.
- E. All plug-ins require SNMP trap forwarders.
- F. There is one plug-in for Grid Control.

Answer: ACD

NEW QUESTION 37

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: ADEF

NEW QUESTION 41

You are planning the physical installation of two full rack Database Machines and two full-rack expansion racks. The four racks will be combined into one multi-rack system.

Which are the two guidelines for installing this configuration in your data Center?

- A. All Database Machines must be placed side by side with no space between them.
- B. All Expansion Racks must be placed side by side with no space between them.
- C. All racks must be placed in such a way that the exhaust air of one rack does not enter the air inlet of another
- D. All racks must be placed side by side with no space between them.
- E. All racks must be isolated from each other with at least one meter between them.
- F. Racks may be placed as required in the machine room.
- G. Expansion Racks must be placed side by side at least one meter apart.

Answer: AC

Explanation:

A: Group related racks together – for example, racks that run a common database or are part of a common cluster

C: Inadequate cold air flow could result in higher air inlet temperatures in the servers due to exhaust air recirculation

NEW QUESTION 42

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: ABC

NEW QUESTION 43

Identify two permitted uses of external InfiniBand connections to a Database Machine.

- A. To connect an ExaLogic Elastic Cloud Machine
- B. To monitor the InfiniBand network using Enterprise Manager
- C. To connect an external tape library

- D. To use a bonded client access network
- E. To use a bonded management network

Answer: AC

Explanation:

A: Combining the Oracle Exalogic Elastic Cloud with the Oracle Exadata Database Machine for SAP NetWeaver The InfiniBand fabric that spans Exalogic and Exadata components provides the following key ways of simplifying and accelerating SAP NetWeaver installations running on Exalogic

Note: A high specialized database networking protocols connects all the components inside an Exadata Database Machine. External connectivity to the Exadata Database Machine is provided through standard 1 Gigabit and 10 Gigab Multiple X3 even larger configurations.

NEW QUESTION 47

You have a partitioned database grid on an X3-2 full rack with two four-node RAC clusters called CLUSA and CLUSB. The storage grid, however, has not been partitioned.

Which files on which servers must be modified after connecting an Exadata storage full expansion rack to your X3-2 Exadata Database Machine on the InfiniBand network so that the cells on the expansion rack are added to the storage grid?

- A. The CELLINIT.ORA files on database servers in CLUSA
- B. The CELLIP.ORA files on the database servers in CLUSA
- C. The CELLINIT.ORA files on the database servers in CLUSB
- D. The CELLIP.ORA files on all existing and newly added Exadata storage servers.
- E. The CELLIP.ORA files on the database servers in CLUSB

Answer: BE

Explanation:

Note:

* cellinit.ora, cellip.ora

-- on database server

cellinit.ora - identifies the storage network interface on the database server cat /etc/oracle/cell/network-config/cellinit.ora

cellip.ora - identifies the Exadata cells that are accessible to the database server cat /etc/oracle/cell/network-config/cellip.ora

* The cellip.ora is the configuration file, on every compute node, that tells ASM instances which cells are available to this cluster.

Here is a content of a typical cellip.ora file for a quarter rack system:

```
$ cat /etc/oracle/cell/network-config/cellip.ora cell="192.168.10.3"
```

```
cell="192.168.10.4" cell="192.168.10.5"
```

NEW QUESTION 51

You are about to run the oplan utility to patch the servers on your test Database Machine before patching the production environment.

The following task might be performed:

- A) Test the failback procedure
 - B) Run the exachk utility
 - C) Read the README file.
 - D) Automate the patch application process as appropriate.
 - E) Verify that the patch provides the functionality it is meant to.
 - F) Apply the patch.
 - G) Evaluate the system performance.
- In which order should the tasks be performed to patch in the recommended fashion?

- A. C, B, D, F, B, E, A
- B. C, D, F, B, E, G, A, B
- C. C, B, D, F, E, G, A
- D. C, B, D, F, E, A, G
- E. C, B, D, F, B, E, G, A

Answer: E

NEW QUESTION 56

Which three are true regarding the use of Storage Indexes?

- A. Different storage regions may have different columns indexed for the same table.
- B. A Storage index is automatically maintained by CELLSRV based on the filter columns of the offload SQL.
- C. The use of Storage indexes for a particular database can be disabled by using an I/O Resource Manager Database Plan.
- D. Storage Indexes occupy space in the Smart Flash Cache.
- E. The use of Storage Indexes for particular categories of I/O can be disabled by using an I/O Resource Manager Category Plan.
- F. A maximum of eight table columns for any table are Indexed per storage region.

Answer: ABF

Explanation:

Note:

* Storage indexes are used during smart scans. All the limitations to smart scans apply to storage indexes. They do not work with joins. Bind variables are supported, however it's slightly more restrictive than regular indexes/queries.

* The storage index is stored in the memory on each of the Exadata storage cells and is created and maintained transparently. However, 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.

* Storage Indexes are a very powerful capability provided in Exadata storage that helps avoid I/O operations. The Exadata Storage Server Software creates and maintains a Storage Index (that is, metadata about the database objects) in the Exadata cell. The Storage Index keeps track of minimum and maximum values of columns for tables stored on that cell. When a query specifies a WHERE clause, but before any I/O is done, the Exadata software examines the Storage Index to determine if rows with the specified column value exist in the cell by comparing the column value to the minimum and maximum values maintained in the Storage

Index. If the column value is outside the minimum and maximum range, scan I/O for that query is avoided. Many SQL Operations run dramatically faster because large numbers of I/O operations are automatically replaced by a few lookups. To minimize operational overhead, Storage Indexes are created and maintained transparently and automatically by the Exadata Storage Server Software.

NEW QUESTION 57

Which three must be true for Smart Scans to be done?

- A. Executing a query in parallel
- B. Setting `_serial_direct_read=true` in the session issuing the SQL statements
- C. Having direct path reads used at run time
- D. Having a 4 meg AU size for the ASM diskgroup containing the tablespace in which tables accessed by a query are stored
- E. `Cell_offload_process = true` for the ASM diskgroup containing the tablespace in which tables accessed by a query are stored.
- F. `cell.smart_scan_capable=true` for the ASM diskgroup containing the tablespace in which tables accessed by a query are stored.

Answer: BCF

NEW QUESTION 61

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.

Answer: ACD

Explanation:

* 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.

NEW QUESTION 66

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 of 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

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 (IORM) 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.

NEW QUESTION 67

Which two communication methods are used by which components in the Enterprise manager Architecture for the Database Machine?

- A. SNMP traps for alerts are sent by the storage server ILOM to the storage server MS process
- B. SNMP traps for alerts are sent by the storage server MS process to the storage server ILOM
- C. SNMP traps for alerts are sent by the storage server ILOM to the Enterprise Manager agent.
- D. SNMP traps for alerts are sent by the storage server MS process to the enterprise Manager agent
- E. SNMP traps for alerts are sent by the storage server ILOM to the storage server RS process.

Answer: AD

Explanation:

There are two types of server alerts that come from Oracle Exadata Storage Server:

* (A) For Integrated Lights Out Manager (ILOM)-monitored hardware components, ILOM reports a failure or threshold exceeded condition as an SNMP trap, which is received by MS.

MS processes the trap, creates an alert for the storage server, and delivers the alert via SNMP to Oracle Enterprise Manager 12c.

* (D) For MS-monitored hardware and software components, MS processes a failure or threshold exceeded condition for these components, creates an alert, and delivers the alert via SNMP to Oracle Enterprise Manager Cloud Control 12c.

Reference: Managing Oracle Exadata with Oracle Enterprise Manager 12c, Oracle White Paper

NEW QUESTION 72

Which two statements are true about the IPTables firewall configuration on a Database Machine- Machine after the default Initial deployment?

- A. IPTables is configured with Oracle supplied rules on the cells.
- B. IPTables is configured with Oracle supplied rules on the database servers.
- C. IPTables is installed and available but not configured on any servers.
- D. IPTables is installed and available but not configured on the database servers.
- E. IPTables is installed and available but not configured on the cells.

Answer: AD

NEW QUESTION 77

Which two are true concerning the allocation of I/O resources by the IORM within the CELLSRV process?

- A. Control File I/O is managed automatically at high priority by IORM.
- B. Control File I/O is considered part of the SYSTEM resource Consume group by IORM.
- C. Log Writer I/O to the Smart Flash Log is considered part of the SYSTEM resource consumer group by IORM.
- D. Log Writer I/O to the Smart Flash Log is managed automatically at high priority by IORM.
- E. Database Writer I/O is managed automatically at normal priority by IORM.
- F. Database Writer I/O is considered part of the SYSTEM resource consumer group by IORM.

Answer: AE

Explanation:

Note:

* IORM Rules

IORM is only "engaged" when needed.

/ (A) Redo and control file writes always take precedence.

/ (E) DBWR (database writer) writes are scheduled at the same priority as user IO.

/ IORM does not intervene if there is only one active consumer group on one database.

/ Any disk allocation that is not fully utilized is made available to other workloads in relation to the configured resource plans.

/ Background IO is scheduled based on their priority relative to user IO.

/ For each cell disk, each database accessing the cell has one IO queue per consumer group and three background queues.

/ Background IO queues are mapped to "high", "medium", and "low" priority requests with different IO types mapped to each queue.

/ If no intradatabase plan is set, all non-background IO requests are grouped into a single consumer group called OTHER_GROUPS.

Reference: Using IORM with Exadata

NEW QUESTION 80

You have altered an index supporting a constraint to be invisible on a large read only data warehouse table, to determine if Smart Scan operations will be fast enough to satisfy your performance requirements.

Given the results of your testing, you consider dropping the index.

Which two statements are true?

- A. You must retain the index and set the constraint to DISABLE NOVALIDATE RELY to enforce the constraint
- B. You may drop the index and use a constraint with the DISABLE NOVALIDATE RELY flags
- C. You must retain the index and make it visible again for the constraint to be enforced.
- D. You may drop the index and make the constraint invisible, because this is enough for the constraint to be enforced.
- E. You may retain the index, and leave it as invisible, because this is enough for the constraint to be recognized.

Answer: AC

Explanation:

Note:

* You may have noticed that we introduced Invisible Indexes as an 11g New Feature. Their main benefit is that we can test whether performance differs if we would drop an index without actually dropping it. This is particular useful after an Exadata Migration because we expect that some conventional indexes migrated are now obsolete and may be substituted by Storage Indexes.

* With making indexes invisible, we can easily check whether indexes are useful without having to drop (and in case recreate) them actually. While this may be of interest for "ordinary" Oracle Databases already, it is particular a useful feature for Exadata where we expect some conventional indexes to become obsolete after a migration.

* DISABLE NOVALIDATE RELY means: "I don't want an index and constaraint checking to slow down my batch data loading into datawarehouse, but the optimizer can RELY on my data loading routine and assume this constraint is enforced by other mechanism". This information can greatly help optimizer to use correct materialized view when rewriting queries. So if you don't use materialized views for query rewrite then you can put RELY for all your constraints (or NORELY for all your constraints) and forget about it.

NEW QUESTION 85

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 1Z0-027 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 1Z0-027 Product From:

<https://www.2passeasy.com/dumps/1Z0-027/>

Money Back Guarantee

1Z0-027 Practice Exam Features:

- * 1Z0-027 Questions and Answers Updated Frequently
- * 1Z0-027 Practice Questions Verified by Expert Senior Certified Staff
- * 1Z0-027 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * 1Z0-027 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year