



Oracle

Exam Questions 1Z0-027

Oracle Exadata Database Machine Administration, Software Release 11.x

NEW QUESTION 1

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 2

You are examining the existing IORM configuration on the cells of Database Machine, to see if they require any 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:          dmor1cel04_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 3

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 4

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 5

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: AC

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.

NEW QUESTION 6

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 7

Which three are true about Exadata storage server alerts?

- A. A threshold based alert gets cleared automatically when the measured value no longer violates the threshold.
- B. A storage server alert is only ever issued as a warning or at a critical situation.
- C. Storage server alerts are all stateless alerts.
- D. Storage server alerts notifications may be sent using SNMP.
- E. Storage server alerts are all stateful alerts.
- F. Storage server alerts notifications may be sent using SMTP.

Answer: ABD

Explanation:

Incorrect:
Not C, Not E: there are both stateful and stateless alerts

NEW QUESTION 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: 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 9

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.

NEW QUESTION 10

You must apply patches and patch bundles in a rolling fashion, if possible, on the components of your Database Machine. You use RAC for your database and also use Data Guard, having standby database on another Database Machine. You wish to have scripts that contain the appropriate command to patch your environment. Your patch bundle is on the first database server and is located at /u01/stage. You have downloaded the oplan utility to the first database server and run the following:
\$ORACLE_HOME/oplan/oplangenerateApplySteps /u01/stage Which two are true concerning oplan?

- A. It will generate instructions for patching only storage servers.
- B. It will generate instructions for patching all the components on the Database Machine.
- C. It will generate instructions for patching the RAC Oracle Home on the primary location.
- D. It will generate instructions for patching the Data Guard environment on the standby location.
- E. It will generate instructions for patching the Grid Infrastructure home on the primary location.

Answer: CE

Explanation:

Oplan generates instructions for all of the nodes in the cluster.

Note:

* Oplan is a utility that facilitates you with the application of bundle patches on Exadata compute nodes via Opatch. This new utility helps you with the patch process by generating step-by-step instructions telling you how to apply a bundle patch in your environment.

* Exadata is much more than a "database-in-a-box." It is a set of compute nodes (think RAC node servers) combined with ultra-fast infiniband (...and 10GB ethernet, and multiple

* 1GB ethernet) and storage nodes.

Incorrect:

Not D: There is no support for Oracle DataGuard

NEW QUESTION 10

You are in the planning stage of the network configuration for your Database Machine. The requirements are:

1. A fault-tolerant network, providing higher availability for connections to database instances
2. Fault tolerance providing higher availability for connections to perform management functions on the database and storage servers.
3. Full monitoring of all Database Machine components using Enterprise Manager

Which three components require external Ethernet network cables to connect your existing network infrastructure to your database machine to satisfy this requirement?

- A. Database servers
- B. Exadata storage servers
- C. InfiniBand storage servers.
- D. Power distribution units (PDUs)
- E. Cisco Ethernet switch

Answer: ADE

NEW QUESTION 15

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 19

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 20

Which statement is true about operating systems on database Machine multi-rack configuration consisting of two full racks and one Exadata storage rack?

- A. All Exadata storage servers and all database servers must run the Oracle Linux O/S and X3-8 database server may run the Oracle Solaris O/S.
- B. All Exadata storage servers must run the Oracle Linux O/S and all database servers within the same cluster must run the Oracle Linux O/s.
- C. All Exadata storage servers must run the Oracle Linux O/S and all database servers within the same cluster must run the same O/s.
- D. All Exadata Storage Servers must run the Oracle Solaris O/S and all database servers within the same cluster must run the same O/S.
- E. All Exadata storage servers in the same cluster must run the same O/S but Exadata Storage Servers in different clusters may run a different O/S.

Answer: C

NEW QUESTION 22

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: AFG

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

NEW QUESTION 26

Identify three valid configuration options that can be implemented during the initial configuration process of a new Database Machine by using the Exadata Deployment Assistant

- A. database character set
- B. size of the RECO ASM diskgroup
- C. size of the DBFS diskgroup
- D. O/S owner mode (standard O/S authentication or role-separated O/S authentication)
- E. RDBMS DBA group name

Answer: BDE

Explanation:

D: Determine which OS owner mode to use

NEW QUESTION 28

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 33

Which two are true about the use of DBFS in a Database Machine environment?

- A. DBFS must be used to bulk load data into a database on the Database Machine if the staging area requires Exadata based shared storage.
- B. DBFS must be used to have a POSIX compliant shared storage solution that is accessible from the database servers on a Database Machine.
- C. DBFS must be used to bulk load data into a production database on the Database Machine.
- D. DBFS must use the DBFS_DG diskgroup for any DBFS store.
- E. DBFS must be used to have a POSIX-compliant Exadata-based shared storage solution.

Answer: CD

NEW QUESTION 36

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 41

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 45

Which two are true about Smart Scan?

- A. a query rewrite may occur to a container table stored in Exadata but will never benefit From Smart scan.
- B. Column projection does not contribute to the performance benefit of Smart Scan
- C. It is possible to offload single row functions to the storage servers.
- D. Some joins can be offloaded to the storage servers.
- E. A query rewrite may occur to a container table stored Exadata, and it will always benefit from Smart Scan.
- F. All joins can be offloaded to the storage servers.

Answer: CD

Explanation:

C: 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.

D (not F):

* Exadata performs joins between large tables and small lookup tables, a very common scenario for data warehouses with star schemas. Joining large tables and small lookup tables is implemented using Bloom Filters, which are a very efficient probabilistic method to determine whether a row is a member of the desired result set.

* If storage indexes are so great, why doesn't Oracle Exadata use them all the time? The short answer is that they are created and used only when they will be beneficial.

* To use storage indexes, Oracle Exadata queries must use smart scans, so not all types of applications can benefit from storage indexes. Applications with queries that include predicates and perform a lot of full table scans or fast full scans of indexes—typically those used in data warehousing environments—will benefit greatly from storage indexes. Online

* transaction processing (OLTP) applications, on the other hand, typically access a small number of rows through standard indexes and do not perform full table scans, so they may not benefit from storage indexes.

Note:

* Storage indexes reside in the memory of the storage servers—also called storage cells—and significantly reduce unnecessary I/O by excluding irrelevant database blocks in the storage cells.

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

Incorrect:

Not B: 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.

Reference: Oracle Communications Data Model Implementation and Operations Guide, Exadata Smart Scan Processing and Storage Index

NEW QUESTION 46

You configuring has two half racks, one with high capacity disks and other with high performance disks and high capacity expansion half rack.

There are two RAC clusters, one on each half rack, which have separate storage grids, each consisting of some of the storage servers in the configuration.

You are planning your deployment of Enterprise Manager to monitor all the components of this multi-rack Database Machine, and must provide for high availability of the monitoring infrastructure.

If the host running the agent which has database machine targets bound to it fails, the monitoring of these targets must be another agent.

Which two are true regarding the configuration used to support this?

- A. Enterprise manager support must be deployed to only one Enterprise Manager Agent in each cluster.
- B. A secondary agent must be deployed on a database server in the same cluster as the server hosting the primary agent.
- C. Enterprise Manager support must be deployed to all Enterprise Manager Agents in each cluster.
- D. A secondary agent may be deployed on a database server in a different cluster than the server hosting the primary agent.
- E. A secondary agent must be deployed on a database server in a different cluster than the server hosting the primary agent.
- F. Enterprise Manager support must be deployed to at least two Enterprise Manager Agents in each of the two RAC clusters.

Answer: AD

Explanation:

Note:

* Instructions for configuring a high availability solution for the Exadata Storage cell or any other Exadata plug-in are documented in the Oracle Database Machine Monitoring Best Practices (Doc ID 1110675.1) document located in My Oracle Support

* High Availability for Plug-Ins

Normally a plug-in target is bound to a specific agent. If the agent is down, the target cannot be monitored.

A procedure exists to facilitate target failover to a secondary agent.

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

* The Oracle ILOM plug-in monitors the Oracle ILOM service processor in a compute node for hardware events and records sensor data to the Oracle Enterprise Manager Repository.

The ILOM plug-in is deployed to the Enterprise Manager (EM) Agent on the first compute node in an Oracle Database Machine, and only that EM agent communicates with the EM Management Server and Repository for all ILOM database server service processors in the Oracle Database Machine.

* OEM Agent and Exadata Plug-ins

1. OEM Agent and Plug-Ins are deployed under Oracle Home
 2. Agent incorporates additional functionality in Plug-In
 3. Agent and Plug-in
- DB Server on Exadata DBM
3. Agent and Plug-in communicates with Storage Server
 4. Plug-Ins are available for monitoring all Exadata hardware components
 5. OEM Agent communicates with the OEM

NEW QUESTION 49

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: ABD

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 time. This algorithm will significantly smooth out redo write response times and provide overall better database performance.

D: Category plans are configured and enabled using the CellCLI utility on the cell. Only one category plan can be enabled at a time.

Incorrect:

Not C: By default, 512 MB of the Exadata flash is allocated to Smart Flash Logging.

NEW QUESTION 54

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 58

You are monitoring and evaluating a create index statement on your Database Machine and have run the following query after executing statement, providing the output. Shown:

```
SQL> SELECT DISTINCT event, total_waits, time_waited/100 wait_secs,
2 average_wait/100 avg_wait_secs
3 FROM V$SESSION_EVENT e, V$MYSTAT s
4 WHERE event LIKE 'cell%' AND e.sid = s.sid;
```

EVENT	TOTAL_WAITS	WAIT_SECS	AVG_WAIT_SECS
cell list of blocks physical read	1	0	.0006
cell single block physical read	1349704	683.94	.0005
cell smart table scan	9191	3.29	.0004

Select two reasons why the statement would have produced so many “cell single block physical read” waits compared to “cell smart table scan” waits.

- A. There are huge numbers of migrated rows in the table on which the index is being built.
- B. There is an uncommitted transaction that has modified one block of the table on which the index is being built, in each cell.
- C. There is a transaction that has modified one block of the table on which the index is being built in each cell, which committed after the create index began.
- D. There are huge numbers of chained rows in the table on which the index is being built.
- E. There is a ROWID column in the table on which the index is being built.

Answer: AD

Explanation:

- A: It could be that row migration.
- D: It could be that row migration or chained rows could cause it.

Note:

* Some facts about scans: Scans exists in “OLTP” systems

Exadata smart scan requires a direct path read.

A direct path read is chosen at runtime based on internal heuristics The STORAGE clause in an explain plan doesnt necessarily mean you will perform a smart scan.

* The buffer caching in certain “OLTP” environments can occasionally induce conventional reads when smart scan is faster.

* Typically see cell multiblock physical read instead of cell smart table scan waits

* No one-size-fits-all solution can be given here but it is very fixable.

NEW QUESTION 61

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 65

Which three are among the software components that constitute the QoS Management framework?

- A. Cluster Health Monitor (CHM)
- B. Cluster Verification Utility (CLUVFY)
- C. O/S Resource Manager
- D. Grid Infrastructure for a standalone server
- E. Grid Infrastructure for a cluster
- F. OC4J cluster resource

Answer: AEF

Explanation:

A: In conjunction with Cluster Health Monitor, QoS Management’s Memory Guard detects nodes that are at risk of failure due to memory over-commitment. It responds by automatically preventing new connections thus preserving existing workloads and restores connectivity once the sufficient memory is again available.

F: The ora.oc4j is for the QoS (Quality of Service Management), which is only available on Exadata.

Incorrect:

Not B: The Cluster Verification Utility (CVU) performs system checks in preparation for installation, patch updates, or other system changes. Using CVU ensures that you have completed the required system configuration and preinstallation steps so that your Oracle grid infrastructure or Oracle Real Application Clusters (Oracle RAC) installation, update, or patch operation completes successfully.

NEW QUESTION 69

You installed ASR Manager on a stand-alone server and configured Auto Service Request (ASR) for your Database machine and its assets. Which three statements are true about this configuration?

- A. When a component fault occurs, fault telemetry is securely transmitted to Oracle via Simple Network Management Protocol (SNMP).

- B. Simple Network Management Protocol (SNMP) traps are used to send notifications from storage servers to ASR Manager.
- C. When a component fault occurs, fault telemetry is securely transmitted to Oracle via HTTPS.
- D. Simple network Management Protocol (SNMP) traps are used to send notification from Enterprise manager to ASR Manager.
- E. Simple Network Management Protocol (SNMP) traps are used to send notifications from database servers to ASR Manager.
- F. Simple Network Management Protocol (SNMP) traps received by ASR Manager are forwarded to Enterprise Manager.

Answer: ABE

Explanation:

To configure fault telemetry, choose one of the following three options:

- Add SNMP Trap Destinations Using OneCommand (recommended for new installations)
- Add SNMP Trap Destinations for Multiple Servers Using the dcli Utility
- Add SNMP Trap Destinations for a Single Server

NEW QUESTION 71

Which three statements are true about the initial storage configuration after the standard deployment of a new database Machine?

- A. The Data_<DBM_Name> and RECO_<DBM_Name> ASM diskgroups are built on part of the interleaved griddisks.
- B. The Data_<DBM_Name> and RECO_<DBM_Name> ASM diskgroups are built on top of the non-interleaved griddisks.
- C. There is a free space available on the hard disks inside the database servers for possible extension of the /u01 file system.
- D. There is no free disk space available on the hard disk inside the database servers for possible extensions of the /u01 file system.
- E. There is free space available on the flashdisk inside the Exadata storage servers for possible use as flash-based griddisks.
- F. There is no free disk available on the flashdisk inside the Exadata storage servers for possible use as flash-based griddisks.

Answer: BCF

Explanation:

B (not A): Griddisks are the fourth layer of abstraction, and they will be the Candidate Disks to build your ASM diskgroups from. By default (interleaving=none on the Celldisk layer), the first Griddisk that is created upon a Celldisk is placed on the outer sectors of the underlying Harddisk. It will have the best performance therefore. If we follow the recommendations, we will create 3 Diskgroups upon our Griddisks: DATA, RECO and SYSTEMDG.

Note:

- * non-root partition (/u01).
- * Oracle Database files: DATA disk group.
- * Flashback log files, archived redo files, and backup files: RECO disk group.

NEW QUESTION 74

You recently upgraded your Exadata image to the latest release; previously you were using 11.2.0.3.

At the same time, you decide to address some performance problems as follows:

You noticed increased latency for the database log writer, especially during the quarterly battery learn cycle on the cells.

You have complaints of erratic performance from certain write-intensive applications. Which two actions could improve performance in these areas?

- A. Enable write-back flashcache by setting lunWriteCacheMode to Write Back Mode.
- B. Use ALTER TABLE in the database to set CELL_FLASH_CACHE = KEEP for the tables belonging to the affected application.
- C. Configure Smart Flash Log on the cells to use some of these of the space on the cell flash devices.
- D. Configure the table belonging to the affected application using CELLCLI, to the set CELL_FLASH_CACHE = KEEP.
- E. Configure Smart Flash Log on the database server to use server flash memory.

Answer: BC

Explanation:

B: The following command could be used to pin the table CUSTOMERS in Exadata Smart Flash Cache

```
ALTER TABLE customers STORAGE (CELL_FLASH_CACHE KEEP)
```

C: Creating Flash Disks Out Of The Flash Cache

When an Exadata cell is installed, by default, all the flash is assigned to be used as flash cache and user data is automatically cached using the default caching behavior. Optionally, a portion of the cache can be reserved and used as logical flash disks. These flash disks are treated like any Exadata cell disk in the Exadata cell except they actually reside and are stored as non-volatile disks in the cache.

Note:

* Pinning Objects In The Flash Cache

Preferential treatment over which database objects are cached is also provided with the Exadata Smart Flash Cache. For example, objects can be pinned in the cache and always be cached, or an object can be identified as one which should never be cached. This control is provided by the new storage clause attribute, CELL_FLASH_CACHE, which can be assigned to a database table, index, partition and LOB column

* There are two techniques provided to manually use and manage the cache. The first enables the pinning of objects in the flash cache. The second supports the creation of logical disks out of the flash for the permanent placement of objects on flash disks.

NEW QUESTION 78

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 constraits) and forget about it.

NEW QUESTION 81

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

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 First Try
- * 1Z0-027 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The 1Z0-027 Practice Test Here](#)