

100% Money Back Guarantee

Vendor: Oracle

Exam Code: 1Z0-027

Exam Name: Oracle Exadata X3 Administration

Version: Demo

Question No : 1

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: A,D

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 No : 2

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.

Answer: B,E

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.

Question No : 3

Identify 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: B

Explanation:

Note:

*Oracle Exadata QoS Management provides the following benefits:

- Reduces the time and expertise requirements for system administrators who manage Oracle

Real Application Clusters (Oracle RAC) resources

- Helps reduce the number of performance outages
- Reduces the time needed to resolve problems that limit or decrease the performance of your

applications

- Provides stability to the system as the workloads change
- Makes the addition or removal of servers transparent to applications
- Reduces the impact on the system caused by server failures
- Helps ensure that service-level agreements (SLAs) are met
- Enables more effective sharing of hardware resources
- Protects existing workloads from over committed memory-induced server failures

-
- Exadata Storage Virtualization
 - Exadata provides a rich set of sophisticated and powerful storage management virtualization capabilities that leverage the strengths of the Oracle Database, the Exadata software, and Exadata hardware.

*When QoS Management is enabled, individual Exadata nodes are protected from memory related failures. The Memory Guard features tracks real-time memory use and should it detect a node has over-committed memory, will prevent new database requests from being sent until the current load is relieved

Question No : 4

Which two are Oracle recommendations for media based backups performed for a database running on a Database Machine?

- A. Allocate equivalent number of channels and instances per tape drive.
- B. Perform periodic level 0 backups and daily cumulative level-1 backups.
- C. Use the InfiniBand network between the database server and media server.
- D. Configure Recovery Manager (RMAN) channels to connect to the least loaded instances.
- E. Use InfiniBand network between the media server and the storage servers.

Answer: A,C

Explanation:

A:Configure one RMAN channel per tape drive and add tape drives to scale backup rates.

C:Configure the Preferred Network Interface (PNI) to direct the Oracle Secure Backup traffic over the InfiniBand network interface.

Example:

```
ob> lspni (List Preferred Network Interface)
```

```
mediaserver1:
```

```
PNI 1:
```

```
interface: mediaserver1-ib
```

```
clients: dbnode1, dbnode2, dbnode3, dbnode4,  
dbnode5, dbnode6, dbnode7, dbnode8
```

PNI 2:

interface: mediaserver1

clients: adminserver

dbnode1:

PNI 1:

interface: dbnode1-ib

clients: mediaserver1

Note:

*Using the Sun ZFS Backup Appliance as an Oracle RMAN backup target for an Oracle Exadata system delivers much faster backup and recovery, enabling organizations to achieve shorter recovery time objectives and shrink backup windows. The appliance is designed for high sustained read and write I/O performance, and it is connected to the Oracle Exadata system via a high-throughput InfiniBand network fabric.

*As the only unified storage vendor to support InfiniBand as a storage network for backup and restore operations, Oracle is leading the way with native high-bandwidth interconnects.

*The InfiniBand network provides 40 Gb of bandwidth per port between the database servers, storage cells, and the Sun ZFS Backup Appliance. Backup and restore operations can be automatically parallelized across all database nodes, Oracle Exadata storage cells, Sun ZFS Backup Appliance channels, and controllers.

Question No : 5

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: C

Explanation: The network ports can be connected directly between the components or the ports can be bonded together to form a logical interface. The bonded ports can be configured for load balancing, fault

tolerance, and better network utilization

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

Question No : 6

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

Question No : 7

Identify two valid reasons for creating multiple griddisks on a Single cell disk.

- 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: B,D

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 No : 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: B,C

Explanation:

Note:

* physical migration techniques (e.g. Data Guard Physical Standby, Transportable Database, Transportable Tablespaces(B), Rapid Clone) vs. logical migration techniques(C)(e.g. Oracle Data Pump).

*You can use either Data Pump export/import or original export/import for Streams instantiations.

*The Oracle Exadata Database Machine runs the standard Oracle Database. Therefore, any application that runs with the Oracle Database today can be seamlessly migrated to the Exadata Database Machine with no changes to the application.

*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.

Question No : 9

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: D,F

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.

Question No : 10

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: A,B,D

Explanation: The default configuration on Exadata is to have 3 diskgroups - DATA, RECO,

and DBFS_DG.

Question No : 11

Yesterday, an Exadata storage server flashdisk entered the poor performance state.

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

Identify the six steps you must perform to replace flashdisk.

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

Answer: A,D,E,F,H,I

Explanation:

Note:

*The default way to deal with the Flash Storage is to use it completely as Flash Cache. You may think of Flash Cache as a prolongation of the Database Buffer Cache. It is populated automatically by the system with objects deemed useful to cache them.

*The second possibility to deal with the Flash Storage is to take a part of it for building ASM diskgroups upon. All files on these ASM diskgroups will then reside permanently on Flash Storage:

```
CellCLI> drop flashcache
```

```
Flash cache exa5cel01_FLASHCACHE successfully dropped
```

```
CellCLI> create flashcache all size=100g
```

Flash cache exa5cel01_FLASHCACHE successfully created

CellCLI> create griddisk all flashdisk prefix=flashdrive

Question No : 12

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: B,C

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.

Question No : 13

Which two statements are true about the use of direct path loads when selecting from external tables in a database on a Database Machine?

- A.** INSERT INTO . . . SELECT FROM statements, executed serially, which select from external tables, require the APPEND hint to use direct path loading.
- B.** CREATE TABLE . . . AS SELECT statements, which select from external tables, attempt to use in direct path loading automatically.
- C.** CREATE TABLE . . . AS SELECT statements, which select from external tables, require the APPEND hint to use direct path loading.
- D.** INSERT INTO . . . SELECT FROM statements, executed serially, which select from

external tables, are unable to use direct path loading.

Answer: A,B

Explanation: A CTAS(Create table as select)will always use direct path(B, not C)load but IAS(Insert as select)statement will not. In order to achieve direct path load with an IAS statement you must add the APPEND hint to the command(A, not D).

Direct path loads can also run in parallel. You can set the parallel degree for a direct path load

either by adding the PARALLEL hint to the CTAS or IAS statement or by setting the PARALLEL clause on both the external table and the table into which the data will be loaded.

Once the parallel degree has been set at CTAS will automatically do direct path load in parallel

but an IAS will not. In order to enable an IAS to do direct path load in parallel you must alter the session to enable parallel DML.

Note:

*Parallel Direct Path Load

The key to good load performance is to use direct path loads wherever possible. A direct path

load parses the input data according to the description given in the external table definition, converts the data for each input field to its corresponding Oracle data type, then builds a column

array structure for the data. These column array structures are used to format Oracle data blocks

and build index keys. The newly formatted database blocks are then written directly to the database, bypassing the standard SQL processing engine and the database buffer cache.

Reference:Best Practices for Implementing a Data Warehouse on the Oracle Exadata Database Machine;Using CTAS & Exchange Partition Replace IAS for Copying Partition on Exadata

Question No : 14

Your customer wants you to partition the database and storage grids in his X3-2 full rack,

creating database clusters and two storage grids.

One cluster will be used for production and should consist of 6 database servers and 11 cells from the first storage grid.

The other cluster will be used for test and development, and should consist of 2 database servers and 3 cells from the second storage grid.

The storage must be partitioned so that the cells are visible only to the appropriate database servers based on the description above.

What must be done to achieve this?

- A. Configure Exadata realms using Oracle ASM scoped security mode.
- B. Configure Exadata realms using Database scoped security mode.
- C. Edit the CELLIP.ORA file on each database server to contain IP addresses of cells in the storage grid associated with cluster to which that database server belongs.
- D. Edit the CELLINIT.ORA file on each database server to contain IP addresses of cells in the storage grid associated with the cluster to which that database server belongs.
- E. Edit the CELLIP.ORA file on each database server to contain IP addresses of database servers which are allowed access to specific cells in the same storage grid.
- F. Edit the CELLIP.ORA file on each cell to contain IP addresses of database servers in the database server grid that are associated with the storage grid to which that cell belongs.

Answer: C

Explanation: 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"
```

Now that we see what is in the cellip.ora, the grid disk path, in the examples above, should make more sense.

Note:

*cellinit.ora decides which network takes storage traffic.

*cellip.ora - list of cells, new cells can be added dynamically without shutdown

Question No : 15

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: B,C,F

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-rootpartition (/u01).

*Oracle Database files: DATA disk group.

*Flashback log files, archived redo files, and backup files: RECO disk group.

Question No : 16

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: A,B,C

Explanation:

Question No : 17

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

- A. To connect an ExaLogicElastic CloudMachine
- 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: A,B

Explanation: A:Combining the Oracle Exalogic Elastic Cloud withthe Oracle Exadata Database Machine for SAPNetWeaver

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 highly specialized database networking protocol connects all the components inside an Exadata Database Machine. External connectivity to the Exadata Database Machine is provided through standard 1 Gigabit and 10 Gigabit Multiple X3 even larger configurations.

Question No : 18

You are evaluating the performance of a SQL statement that accesses a very large table, and have run the following query producing the output shown:

```
SQL> SELECT s.name, m.value/1024/1024 MB FROM V$SYSSTAT s, V$MYSTAT m
  2  WHERE s.statistic# = m.statistic# AND
  3  (s.name LIKE 'physical%total bytes' OR s.name LIKE 'cell phys%'
  4  OR s.name LIKE 'cell IO%');
```

NAME	MB
physical read total bytes	19047.2266
physical write total bytes	0
cell physical IO interconnect bytes	4808.85828
cell physical IO bytes pushed back due to excessive CPU on cell	0
cell physical IO bytes saved during optimized file creation	0
cell physical IO bytes saved during optimized RMAN file restore	0
cell physical IO bytes eligible for predicate offload	18005.6953
cell physical IO bytes saved by storage index	0
cell physical IO interconnect bytes returned by smart scan	3767.32703
cell IO uncompressed bytes	18005.6953

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 for blocks 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.

Answer: B,E

Explanation:

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 No : 19

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: C,E

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

Question No : 20

Consider the following list of software components:

- 1.DCLI
- 2.Management Server (MS)
- 3.ASM Instance
- 4.RDBMS instance
- 5.Restart Server (RS)
- 6.Cellcli
- 7.Cell Server (CELLSRV)

8.Diskmon

Identify the location where these software components may run in the standard Database machine deployment.

- A. 3, 4 and 8 run on the database servers; 1, 2, 5, 6 and 7 run on the Exadata Storage servers.
- B. 4 and 8 run on the database servers; 1, 2, 3, 5, 6 and 7 run on the Exadata Storage servers.
- C. 1, 3 and 4 run on the database servers; 2, 5, 6, 7 and 8 run on the Exadata Storage servers.
- D. 3, 4 and 8 run on the database servers; 1, 2, 5, 6 and 7 run on the Exadata Storage servers.
- E. 3, 4 and 8 run on the database servers; 1, 2, 5, 7 and 8 run on the Exadata Storage servers.
- F. 1, 3, 4 and 8 run on the database servers; 1, 2, 5, 7 and 8 run on the Exadata Storage servers.

Answer: B

Explanation: * (not C or not F): Exalogic includes a tool called DCLI (Distributed Command Line Interface) that can be used to run the same commands on all or a subset of compute nodes in parallel.

* Cellcli(not E): on Exadata Storage Servers.

* RDBMS instance on database server.

* CELLSRV on Exadata Storage Servers.

* ASM on Exadata Storage Servers(not D)

After an Oracle ASM instance has been installed on a single-instance Oracle Database or in an Oracle Real Application Clusters (Oracle RAC) environment, the final Oracle ASM configuration can be performed.

Question No : 21

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: B,C

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.

Question No : 22

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: A,D

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

Question No : 23

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?
- H. C, B, D, F, B, E, A
 - I. C, D, F, B, E, G, A, B
 - J. C, B, D, F, E, G, A
 - K. C, B, D, F, E, A, G
 - L. C, B, D, F, B, E, G, A

Answer: B

Explanation:

Note:

- *Schedule regular health checks with Exachk
 - Detects bad disks, faulty hardware, incorrect configuration
 - At least every 90 days
 - Before and after patching/upgrades

– Follow recommended best practices

*we are recommending running exachk or Healthcheck (depending on your DBM/Exadata model) *before* maintenance and *after* maintenance.

Question No : 24

Which two may act as building blocks for the creation of two or more griddisks residing on the same physical device and which have more balanced performance characteristics?

- A. griddisk on hard drive-based interleaved celldisks.
- B. griddisk on hard drive-based celldisks consumed by ASM when creating diskgroups using ASM-based intelligent Data Placement (IDP)
- C. griddisk on hard drive-based non-interleaved celldisks.
- D. griddisk on Flash-base celldisks.

Answer: A,B

Explanation: A (not C):Interleaved Griddisks mean equally fast Griddisks, whereas with non Interleaved Griddisks (the default) the Griddisks created first are being taken from the outer tracks of the underlying physical disk.

B: 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.

Question No : 25

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: A

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

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

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

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

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

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

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

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

read and write ranges of bytes to and from disk so when offload processing is not possible

Exadata operates like a traditional storage device for the Oracle Database. But when feasible, the

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

on the Exadata Storage Server so only requested data is returned to the database server.

iDB is built on the industry standard Reliable Datagram Sockets (RDSv3) protocol and runs over

InfiniBand. ZDP (Zero-loss Zero-copy Datagram Protocol), a zero-copy implementation of

RDS, is used to eliminate unnecessary copying of blocks. Multiple network interfaces can

be used on the database servers and Exadata cells. This is an extremely fast low-latency

protocol that minimizes the number of data copies required to service I/O operations.

Note:

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

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

InfiniBand

connectivity for high availability.

*The same InfiniBand network also provides a high performance cluster

interconnect for the Oracle Database Real Application Cluster (RAC) nodes.

Note:

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

Question No : 26

In which two locations should files be staged, to be loaded using external tables into a database on a Database Machine?

- A. On a dbfs file system stored in a staging database on the Database Machine
- B. On an Exadata-based ACFS file system on the Database Machine
- C. On an nfs file system mounted on a database server where the external table will be accessed.
- D. On local storage on one or more cells that are accessible to the database server where the load will be performed.

Answer: C,D

Explanation: Prepare a Temporary File System

The Oracle E-Business Suite Rapid Install creates a file system based Oracle Database release

11.1.0.7 database, and it cannot install directly to a database that uses Oracle ASM.

Therefore it

requires non-Exadata storage of about 55 GB and an additional 50 GB of storage for the install

staging area for a total of 105 GB. The file system to be used for the database node file system

and the stage area can either use the database machine local storage or an NFS mounted file

system.

Reference: Installing Oracle E-Business Suite Release 12 with the Oracle Exadata Database Machine

Question No : 27

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`

Question No : 28

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 logfile completes first.
- C. Smart Flash Log is enabled by default, using 1024 MB of Rash 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 LogFor different I/O categories.

Answer: A,B,D

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

B: Smart Flash Logging works as follows. When receiving a redo log write request, Exadata will do parallel writes to the on-disk redo logs as well as a small amount of space reserved in the

flash

hardware. When either of these writes has successfully completed the database will be immediately notified of completion. If the disk drives hosting the logs experience slow response

times, then the Exadata Smart Flash Cache will provide a faster log write response time.

Conversely, if the Exadata Smart Flash Cache is temporarily experiencing slow response times

(e.g., due to wear leveling algorithms), then the disk drive will provide a faster response 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

Question No : 29

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: B,E

Explanation:

Note:

To Read the [Whole Q&As](#), please purchase the [Complete Version](#) from [Our website](#).

Trying 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 **69,000** Satisfied Customers Worldwide
- ★ Multi-Platform capabilities - **Windows, Mac, Android, iPhone, iPod, iPad, Kindle**

Need Help

Please provide as much detail as possible so we can best assist you.

To update a previously submitted ticket:



 One Year Free Update <p>Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.</p>	 Money Back Guarantee <p>To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.</p>	 Security & Privacy <p>We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.</p>
---	---	--

[Guarantee & Policy](#) | [Privacy & Policy](#) | [Terms & Conditions](#)

Any charges made through this site will appear as Global Simulators Limited.

All trademarks are the property of their respective owners.

Copyright © 2004-2015, All Rights Reserved.