

Oracle-1z0-027-examcollection-by-Yan-72q-08-05-2019

Free exam dumps provided by IT-Libraries Team :)

Note:

Before attending the exam we strongly recommend studying our premium questions.
In order to maximize your chances of success take a look on our free study guides.

Feedback

We are committed to product quality, if you have any questions about our product please contact us at support@itlibraries.com

Always check our website to make sure you are using the most updated version!

©2019 www.itlibraries.com

Special thanks to Yan for sharing this.

Exam A

QUESTION 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

Correct Answer: AD

Explanation

Explanation/Reference:

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 2

Which two are regarding the case of storage indexes?

- A. To increase the chance of using the a storage index, you can make table indexes invisible.
- B. To maximize the benefit of storage Indexes, load your data stored on the filtered columns.
- C. The cell physical 10 bytes saved by storage index statistic returns multiple rows, one for each storage server.
- D. Storage indexes are retained after a cell is rebooted.
- E. Avoid the use of bind variables because Storage Indexes do not work with bind variables.

Correct Answer: AB

Explanation

Explanation/Reference:

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.

QUESTION 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

Correct Answer: A

Explanation**Explanation/Reference:****QUESTION 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.

Correct Answer: AC

Explanation**Explanation/Reference:**

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 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

Correct Answer: B

Explanation

Explanation/Reference:

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

QUESTION 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

Correct Answer: D

Explanation

Explanation/Reference:

Airflow must be front-to-back.

Reference: Oracle White Paper, ORACLE Exadata Database Machine X3-8

QUESTION 7

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

- A. To segregate storage into multiple pools with different performance characteristics
- B. To facilitate normal or high redundancy ASM diskgroups
- C. To enable disk mirroring for the system area
- D. To segregate storage into multiple pools that can be assigned to different databases
- E. To segregate storage into multiple pools that can be assigned to different resource consumer groups in the same database.

Correct Answer: AD

Explanation

Explanation/Reference:

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

Correct Answer: CE

Explanation

Explanation/Reference:

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 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

Correct Answer: DF

Explanation

Explanation/Reference:

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 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

Correct Answer: ABC

Explanation

Explanation/Reference:

QUESTION 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.

Correct Answer: ADEFHI

Explanation

Explanation/Reference:

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 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.

Correct Answer: BC

Explanation

Explanation/Reference:

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 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.

Correct Answer: AB

Explanation

Explanation/Reference:

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 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.

Correct Answer: C

Explanation

Explanation/Reference:

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 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.

Correct Answer: BCF

Explanation

Explanation/Reference:

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.

QUESTION 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

Correct Answer: ADE

Explanation

Explanation/Reference:

QUESTION 17

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

Correct Answer: AC

Explanation

Explanation/Reference:

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 Gigabit Multiple X3 even larger configurations.

QUESTION 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.

Correct Answer: BD

Explanation

Explanation/Reference:

QUESTION 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/oplanguenerateApplySteps /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.

Correct Answer: CE

Explanation

Explanation/Reference:

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 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 9 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.

Correct Answer: D

Explanation

Explanation/Reference:

QUESTION 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

Correct Answer: BC
Explanation

Explanation/Reference:

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 22

You are configuring a Database Machine with 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.

Correct Answer: AD
Explanation

Explanation/Reference:

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 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?

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

Correct Answer: E

Explanation

Explanation/Reference:

QUESTION 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.

Correct Answer: AD

Explanation

Explanation/Reference:

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.

QUESTION 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

Correct Answer: E

Explanation

Explanation/Reference:

QUESTION 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.

Correct Answer: AC

Explanation

Explanation/Reference:

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

QUESTION 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

Correct Answer: B

Explanation

Explanation/Reference:

Gather all diagnostics information

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

QUESTION 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 log file 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 Log For different I/O categories.

Correct Answer: ABD

Explanation

Explanation/Reference:

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 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

Correct Answer: BE

Explanation

Explanation/Reference:

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"
```

QUESTION 30

Which two are true about Exadata storage server alerts?

- A. Metric alerts are never stateful.
- B. Metrics have no thresholds set on them by default.
- C. SNMP alert notifications can be sent to only one destination.
- D. Metric threshold; if set, will persist across storage server reboots.
- E. SMTP alert notifications must be sent to both ASR manager and Enterprise Manager Agents

Correct Answer: DE

Explanation

Explanation/Reference:

Incorrect:

Not A: there are both stateful and stateless alerts.

Not B: Metrics have defaults.

Not C: An SNMP alert can have multiple destinations.

QUESTION 31

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

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

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

- A. Upgrade source database to 11.2.0 and migrate using a physical standby database.
- B. Migrate using Data Pump.
- C. Migrate using GoldenGate.
- D. Migrate using cross platform Transportable Database.
- E. Migrate using ASM online migration.

Correct Answer: AB

Explanation

Explanation/Reference:

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

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

QUESTION 32

Consider the following sequence of CELLCLI commands and SQL statements:

```
CELLCLI> CREATE CELLDISK ALL HARDDISK
CELLCLI> CREATE GRIDDISK ALL HARDDISK PREFIX=DATA01 SIZE=120G
CELLCLI> CREATE GRIDDISK ALL HARDDISK PREFIX=DATA02 SIZE=120G
SQL> CREATE DISKGROUP DATA02 NORMAL REDUNDANCY DISK 'O/*/DATA02*' ;
SQL> CREATE DISKGROUP DATA01 NORMAL REDUNDANCY DISK 'O/*/DATA01*' ;
```

All CELLCLI commands and SQL statements are executed successfully without warnings.

Which statement is true about the I/O performance of the DATA01 ASM diskgroups?

- A. The DATA01 diskgroups has better performance because the DATA01 griddisks were created first.
- B. The DATA02 ASM diskgroup has better performance because the DATA02 diskgroups was created first.
- C. The DATA01 diskgroup has better performance because the DATA01 griddisks are created on interleaved cell disks.
- D. The performance of both diskgroups is balanced because they are created on interleaved griddisks.

Correct Answer: A

Explanation

Explanation/Reference:

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

Incorrect:

Not D: non-interleaved is the default.

QUESTION 33

You must drop all celldisks on all the cells in a quarter rack as part of a reconfiguration project, to support normal redundancy interleaving.

Select two actions that describe the operating system (O/S) account on the cells to which you should log in, and the tool that may be used to drop the celldisks.

- A. To the CELLMONITOR account using CELLCLI interactively on each cell.
- B. To CELLADMIN account calling CELLCLI on all cells using DCLI
- C. To either the CELIMONITOR or the CELLADMIN account calling CELLCLI on all cells using DCLI
- D. To the CELLMONITOR account calling CELLCLI on all cells using DCLI
- E. To the CELLADMIN account using CELLCLI interactively on each cell
- F. To either the CELLMONITOR or CELLADMIN account calling interactively on each cell

Correct Answer: BE

Explanation

Explanation/Reference:

Note:

* Each Exadata Storage cell can be monitored, configured, and maintained using the cellcli command line interface. To invoke the CLI, simply login to the Exadata cell as cellmonitor, celladmin, or root, and type "cellcli". Within cellcli, a wide range of commands can be invoked to monitor the cell.

* While connected to cellcli as the Linux account celladmin, you can perform all cellcli commands except calibrate.

While logged into the cell server as cellmonitor, you can do any "read-only" commands; you cannot change any configuration or attributes.

QUESTION 34

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.

Correct Answer: AD

Explanation

Explanation/Reference:

QUESTION 35

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

Correct Answer: ABC

Explanation

Explanation/Reference:

QUESTION 36

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

- A. ASM intelligent Data Placement (IDP)
- B. Intelligent Database Protocol (IDB)
- C. Database Resource Manager (DBRM)
- D. I/O Resource Manager (IORM)
- E. Database Filesystem (DBFS)
- F. The DISKMON process

Correct Answer: ABD

Explanation

Explanation/Reference:

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.

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

QUESTION 37

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.

Correct Answer: ABE

Explanation

Explanation/Reference:

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

QUESTION 38

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.

Correct Answer: AE

Explanation

Explanation/Reference:

QUESTION 39

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:          dmorlcel04_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 flashdms 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.

Correct Answer: AD

Explanation

Explanation/Reference:

QUESTION 40

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

- A. Decompression can be done on the database servers.
- B. Compression can be done on the Exadata storage servers.
- C. Compression can be done on the database servers.
- D. Decompression can be done on the Exadata storage servers.

Correct Answer: ACD

Explanation

Explanation/Reference:

* decompression

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

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

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

QUESTION 41

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.

Correct Answer: AD

Explanation

Explanation/Reference:

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 doesn’t 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.

QUESTION 42

Which three factors should you consider when choosing a method for migrating a database to the Database Machine?

- A. Type of database workload
- B. Number of tablespaces in the source database
- C. Size of the source database
- D. Endian format of the source database
- E. ASM attributes and allocation unit (AU) size of the diskgroups used by the source database

Correct Answer: ACD

Explanation

Explanation/Reference:

A: Use Real Production Workload

– Real Application Testing (RAT)

• Database Workload Replay

• SQL Performance Analyzer (SPA)

D:

* Endian format: Non-Exadata format

Reference: Exadata MAA Best Practices, Migrating Oracle Databases

QUESTION 43

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

Correct Answer: AEF

Explanation**Explanation/Reference:**

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.

QUESTION 44

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.

Correct Answer: AD

Explanation**Explanation/Reference:**

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

QUESTION 45

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.

Correct Answer: AFG

Explanation

Explanation/Reference:

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

QUESTION 46

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.

Correct Answer: CD

Explanation**Explanation/Reference:**

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

QUESTION 47

A read-only application is in development and is using a test database on a Database Machine. You are examining SQL statements from this application in an attempt to determine which ones will benefit from the Exadata Smart scan capability.

The following is true about the tables used by the application:

1. The data for the tables has just been loaded.
2. There are no applications accessing the tables currently.
3. None of the indexes are compressed or reverse key indexes.
4. The tables use the default organization type data.

5. The only data types for the table are varchar (2), number, or date.
6. The largest number of columns for any table is 29.
7. No access is based on ROWID, or virtual columns.

Which two access paths will always generate either a set of “cell smart table scan” or a set of “cell smart index scan” requested?

- A. Full scans on sorted hash clustered tables executed in parallel
- B. Full table scans on index organized tables executed in parallel.
- C. Full table scans on heap tables executed in parallel
- D. Full scans on index clustered tables executed in parallel
- E. full scans on hash clustered tables executed in serially
- F. fast full scans on B*Tree indexes executed in parallel
- G. full index scans on B*Tree indexes executed in parallel

Correct Answer: CF

Explanation

Explanation/Reference:

QUESTION 48

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

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

- A. The targets must be configured by an Enterprise Manager administrator manually using Grid Control.
- B. Setupem.sh is used again to configure the targets.
- C. The targets must be configured by an O/S administrator manually.
- D. The targets must be configured by a database administrator manually.
- E. There may be more than one target for each plug-in.

Correct Answer: AE

Explanation

Explanation/Reference:

After running the `setupem.sh` script:

- The Cloud Control OMS and Repository will be installed
- The Exadata plug-ins are deployed on the OMS

QUESTION 49

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

Correct Answer: BDE

Explanation

Explanation/Reference:

D: Determine which OS owner mode to use

QUESTION 50

Which two are true about the use of the Integrated Lights Out Manager (ILOM) on the Database Machine?

- A. ILOM can be used to power-on the Cisco switch.
- B. ILOM can be used to power on the InfiniBand switches.
- C. ILOM can be used to power-on the database servers.
- D. ILOM generates hardware alerts for the power distribution units.
- E. ILOM provides a remote console for the storage servers.

Correct Answer: CE

Explanation

Explanation/Reference:

* The Oracle ILOM Remote Console is supported on all Oracle Sun x86 processor-based servers. It is also supported on some SPARC processor-based servers. The Oracle ILOM Remote Console is a Java application that you can launch from the Oracle ILOM web interface

* Remote Management: Integrated Lights Out Manager (ILOM) Ethernet port

QUESTION 51

You are checking the status of the ports on the one of the InfiniBand switches in your Database Machine.

You have run the getportstatus as shown:

```
[root@exampsw-ib2 bin]# getportstatus 17b
Port status for connector 17B Switch Port 2
Adminstate:.....Enabled
LinkWidthEnabled:.....1X or 4X
LinkWidthSupported:.....1X or 4X
LinkWidthActive:.....4X
LinkSpeedSupported:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
LinkState:.....Active
PhysLinkState:.....LinkUp
LinkSpeedActive:.....10.0 Gbps
LinkSpeedEnabled:.....2.5 Gbps or 5.0 Gbps or 10.0 Gbps
```

You have 36 ports to check.

Which two methods would you use to check the status of all 36 ports?

- A. Use Enterprise Manager.
- B. Use the ibqueryerrors.pl script logged into the InfiniBand switch as root.
- C. Use Database Control.
- D. Create user-defined metrics for the InfiniBand switch.

Correct Answer: AB

Explanation

Explanation/Reference:

http://docs.oracle.com/cd/E17559_01/install.111/e20085/e20085-03.htm

QUESTION 52

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

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

- A. Column filtering
- B. Sort merge join filtering
- C. Predicate filtering
- D. Nested loop Join filtering
- E. Hash join filtering

F. Virtual column filtering

Correct Answer: ACE

Explanation

Explanation/Reference:

A: Smart Scan Column Filtering

Exadata provides column filtering, also called column projection, for table scans. Only the columns requested are returned to the database server rather than all columns in a table. For example, when the following SQL is issued, only the employee_name and employee_number columns are returned from Exadata to the database kernel.

```
SELECT employee_name, employee_number FROM employee_table.
```

For tables with many columns, or columns containing LOBs (Large Objects), the I/O bandwidth saved can be very large. Using both predicate and column filtering dramatically improves performance and reduces I/O bandwidth consumption. In addition, column filtering also applies to indexes, allowing for even faster query performance.

C: Smart Scan Predicate Filtering

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

```
SELECT * FROM employee_table WHERE hire_date > '1-Jan-2003'.
```

This ability to return only relevant rows to the server greatly improves database performance. This performance enhancement also applies as queries become more complicated, so the same benefits also apply to complex queries, including those with subqueries.

QUESTION 53

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.

Correct Answer: BCF

Explanation

Explanation/Reference:

QUESTION 54

Identify the correct sequence of commands to completely power off a Database Machine in an Orderly fashion:

1. Execute `crsctl stop cluster -all` as the grid user from one database server.
2. Execute `crsctl stop cluster -all` as the root user from one database server.
3. Power off all network switches using their power switch.
4. Execute `crsctl stop cluster -all` as the root user from one database server.
5. Execute `crsctl stop cluster -all` as the grid user from one database server.
6. Power off the rack using the power switches on the PDUs.
7. Execute `shutdown -h now` on all database servers.
8. Execute `shutdown -h now` on all Exadata storage servers.
9. Execute `shutdown -h now` on all servers.

- A. 1, 9, 3 and 6
- B. 2, 9 and 6
- C. 4, 7, 8, 3 and 6
- D. 5, 8, 7, 3 and 6
- E. 2, 8, 7, 3 and 6
- F. 1, 8, 7, 3 and 6
- G. 2, 7, 8 and 6
- H. 2, 8, 7 and 6

Correct Answer: G

Explanation

Explanation/Reference:

QUESTION 55

You plan to migrate an existing production database supporting online transaction processing (OLTP) workloads to the Exadata Database Machine. The database currently supports an application requiring fast response times, which satisfies stringent requirements, and most of the application queries use index access to the tables in the application schema.

For which case would you consider dropping indexes to allow Smart Scans to occur?

- A. Drop non-constraint indexes if Smart Scan occurs instead of an index access path on the corresponding table.
- B. Drop non-constraint indexes if Smart Scan performs better than index unique scans on the corresponding table.
- C. Drop non-constraint indexes if Smart Scan performs better than index range scans on the corresponding table.
- D. Drop non-constraint Indexes if Smart Scan performs better than any index scans on the corresponding table.

Correct Answer: D

Explanation

Explanation/Reference:

QUESTION 56

Your Database Machine has a large database with some very large tables supporting OLTP workloads.

High volume Insert applications and high volume update workloads access the same tables.

You decide to compress these tables without causing unacceptable performance overheads to the OLTP application.

Which three are true regarding this requirement?

- A. Using 'compress for oltp' will compress the data less than if using Hybrid Columnar Compression when specified with compress for query low.
- B. The compression is performed on the storage servers when using compress for oltp in an Exadata environment.
- C. The compression method compress for archive high is the worst fit for this requirement.
- D. Using 'compress for oltp' will compress the data more than if using Hybrid Columnar Compression when specified with compress for archive low.
- E. The compression is performed on the database servers when using 'compress for oltp' in an Exadata environment.

Correct Answer: ACE

Explanation

Explanation/Reference:

Note:

(E not B):

* Types of compression

Basic compression

OLTP compression

Warehouse compression

Online archival compressio

*

/ OLTP compression allows compression during DML operations.

/ Basic compression works at the data block level.

* When you enable table compression by specifying COMPRESS FOR OLTP, you enable OLTP table compression. Oracle Database compresses data during all DML operations on the table. This form of compression is recommended for OLTP environments.

* When you specify COMPRESS FOR QUERY or COMPRESS FOR ARCHIVE, you enable hybrid columnar compression. With hybrid columnar compression, data can be compressed during bulk load operations. During the load process, data is transformed into a column-oriented format and then compressed. Oracle Database uses a compression algorithm appropriate for the level you specify. In general, the higher the level, the greater the compression ratio.

Hybrid columnar compression can result in higher compression ratios, at a greater CPU cost. Therefore, this form of compression is recommended for data that is not frequently updated.

QUESTION 57

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.

Correct Answer: C

Explanation

Explanation/Reference:

QUESTION 58

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.

Correct Answer: ABF

Explanation

Explanation/Reference:

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.

QUESTION 59

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.

Correct Answer: BCEF

Explanation

Explanation/Reference:

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

QUESTION 60

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.

Correct Answer: AC

Explanation

Explanation/Reference:

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 particularly 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 particularly 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 constraint 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.

QUESTION 61

Which two activities are supported on the storage servers in the Database Machine?

- A. Installing an alternative package manager
- B. configuring secure shell user equivalency for the callmonitor user
- C. changing root password
- D. upgrading the Storage Server software package using RPN
- E. upgrading a device driver for hard disks when inserting a replacement disk after a hard disk failure

Correct Answer: BC

Explanation

Explanation/Reference:

QUESTION 62

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 and 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 to be set because they are preset in the ILOM and these are sufficient for monitoring.

Correct Answer: D

Explanation

Explanation/Reference:

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.

QUESTION 63

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.

Correct Answer: ACD

Explanation

Explanation/Reference:

QUESTION 64

Which is true regarding Enterprise Manager monitoring and configuration?

- A. Agents must be deployed on the storage servers, in case all the database servers are down.
- B. Monitoring the power distribution units (PDUs) is done by Enterprise Manager agents, but additional ethernet cables are required to connect the PDUs to the internal network switch in the Database Machine
- C. Monitoring the power distribution units (PDUs) is done by Enterprise Manager agents without, any additional network configuration.
- D. Agents may be deployed on the storage servers, in case all the database servers are down.
- E. Monitoring the power distribution units (PDUS) is done by Enterprise Manager agents but additional ethernet cables are required to connect the PDUs to the corporate network switch.

Correct Answer: C

Explanation

Explanation/Reference:

Note:

* The following prerequisites must be met before you can deploy the plug-in:

Set and validate the preferred credentials on all Agents where you want to deploy the plug-in. This is essential in order to deploy the plug-in.

The PDU firmware version must be 1.02 or later.

* This plug-in will be used to monitor PDU's actual current value of Phase1, Phase2, and Phase3 and raises appropriate events if they have crossed the different threshold values. Each PDU will have a plug-in instance added to its own OEM Grid Control agent.

By deploying the plug-in to your Grid Control environment, you gain the following management features:

Monitor PDU targets.

Raise alerts and violations based on thresholds set on monitoring and configuration data.

Provide rich out-of-box metrics and reports based on the gather

Reference: Oracle Enterprise Manager System Monitoring Plug-In Installation Guide for Exadata Power Distribution Unit

QUESTION 65

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.

Correct Answer: CD

Explanation

Explanation/Reference:

QUESTION 66

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.

Correct Answer: ABD

Explanation

Explanation/Reference:

Incorrect:

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

QUESTION 67

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

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

- A. 'crsctl stop cluster -all' on one of the database servers
- B. CellCLI> LIST GRIDDISK ATTRIBUTE name WHERE asmdeactivationoutcome != 'Yes'
- C. CellCLI> ALTER GRIDDISK ALL INACTIVE
- D. CellCLI> LIST GRIDDISK WHERE STATUS != 'inactive'
- E. 'shutdown -h now' on the affected cell

Correct Answer: BCDE

Explanation

Explanation/Reference:

B: Step 1:

Run the following command to check if there are other offline disks

CellCLI> LIST GRIDDISK ATTRIBUTES name WHERE asmdeactivationoutcome != 'Yes'

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

C: Step 2:

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

CellCLI> ALTER GRIDDISK ALL INACTIVE

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

D: Step 3:

-- Verify all grid disks are INACTIVE to allow safe storage server shut down by running the following command:
CellCLI> LIST GRIDDISK ATTRIBUTES name, asmmodestatus
CellCLI> LIST GRIDDISK

If all grid disks are INACTIVE, then the storage server can be shutdown without affecting database availability

E: Step 4:
To stop a server, use the shutdown command. To stop immediately and keep it down, i.e. not reboot, execute:

shutdown -h -y now

QUESTION 68

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.

Correct Answer: AE

Explanation

Explanation/Reference:

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

QUESTION 69

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.

Correct Answer: ADEF

Explanation

Explanation/Reference:

QUESTION 70

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

- A. If two consumer groups A and B in the PROD database get 10% and 15% respectively of resource allocation, then the percentage of I/O resource would be the same if they got 20% and 30% respectively, provided that the category plans and inter database plans are unchanged, and consumer groups A and B are still mapped to the same category.
- B. If two consumer groups C and D In the PROD database get 10% and 15% respectively of resource allocation, then the percentage of I/O resource would be the same if they still got 10% and 15% respectively and were remapped to a different category by the DBA, provided that the category plans and Interdatabase plans are unchanged.
- C. If two consumer groups P and Q In the PROD database get 20% and 10% respectively of resource allocation, then the percentage of I/O resource would be the same if they got 10% and 5% respectively, and the Interdatabase plan has changed, provided that the category plan is unchanged, and consumer groups A and B are still mapped to the same categories.
- D. Log Writer I/O and Control File I/O are high priority I/O managed automatically by IORM.
- E. Database Writer I/O is considered to be the highest priority I/O and is managed based on attributes of the IORM plan.

Correct Answer: CD

Explanation**Explanation/Reference:**

*(not A, Not B) Rules in an interdatabase resource plan specify allocations to databases, not consumer groups.

* IORM Rules

IORM is only "engaged" when needed.

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

/ (Not 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

Note:

* Use IORM metrics to track

/ I/O load per Consumer Group (IOPS, MBPS, disk utilization %)

/ I/O throttling per Consumer Group

Reference: Using IORM with Exadata

QUESTION 71

You plan to migrate a database supporting an an OLTP workload to your Database Machine

This is part of a consolidation project and several other databases already exist on the Database Machine.

Which three Exadata features may help to improve the performance of this OLTP workload?

- A. Hybrid Columnar Compression
- B. I/O Resource Manager
- C. Smart Flash Cache
- D. Smart Flash Log
- E. Smart Scan
- F. Storage Index

Correct Answer: BCD

Explanation**Explanation/Reference:**

C:

OLTP performance benefits with Exadata

/ (D) Smart flash log for low latency commits

/ (C) Smart flash cache for low latency reads. KEEP in Flash for critical objects

QUESTION 72

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.

Correct Answer: AC

Explanation

Explanation/Reference:

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