

Exam Questions 1z0-070

Oracle Exadata X5 Administration

<https://www.2passeasy.com/dumps/1z0-070/>



NEW QUESTION 1

Because of a motherboard failure on an X6 high-capacity Exadata storage server, you must replace the server chassis with a replacement unit. Which are three of the components must you move from the original Exadata storage server to the new Exadata storage server?

- A. Redundant power supplies
- B. The host bus channel adaptor
- C. The USB flash drive
- D. The four Sun Flash Accelerator PCIe cards
- E. The 12 physical disks

Answer: ADE

Explanation: Remove the following reusable components: References:https://docs.oracle.com/cd/E62172_01/html/E63689/goiqw.html#scrolltoc

NEW QUESTION 2

Which three statements are true about Exadata storage server alerts in an X5 Database Machine?

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

Answer: ACF

Explanation: Exadata cell (storage server) alerts can be delivered using SMTP or SNMP or both.

Although there are three types of storage alerts: informational, warning and critical, they are issued when threshold metrics reached either to warning or critical.

NEW QUESTION 3

Which two statements are true about RMAN incremental level-1 backups for a database running on an X5 Database Machine?

- A. Block Change tracking (BCT) is not needed for databases in an X5 Database Machine, because ASM automatically does block change tracking when diskgroups are built from griddisks.
- B. cellsrv filters blocks and returns only those that have changed since the last backup, for databases in archive log mode.
- C. cellsrv returns all blocks that have been or might have been changed since the last backup, and RMAN filters the blocks that do not require backup.
- D. cellsrv filters blocks and returns only those that have changed since the last backup, for databases in noarchive log mode.
- E. Block Change Tracking (BCT) is not needed for databases in an X5 Database Machine, because change tracking is done by cellsrv.

Answer: DE

Explanation: D: Whenever data blocks change, the Change Tracking Writer (CTWR) background process tracks the changed blocks in a private area of memory. When RMAN performs an incremental backup on the Exadata platform, cellsrv filters out unwanted blocks and sends back only those that have changed since the last level 0 or level 1 backup.

E: For the Exadata platform, you may choose to allow cellserv to do all the blocking filtering for incremental backups.

Note: RMAN block change tracking is used to improve the performance of incremental backup. The level 0 incremental backup scans the entire database but level 1 incremental backups use the block change tracking file to scan only the blocks that have changed since the last backup. This significantly reduces the amount of reads that are required on the database.

References: Expert Oracle Exadata (29 Aug 2015), page 319

NEW QUESTION 4

You wish to determine if the I/O resource management plan that you created has helped improve the performance of OLTP category I/Os on your X6 Exadata Database Machine.

You decide to examine the relevant metrics on all the cells, to see whether the I/O rate has improved for this category compared to last week, and whether waits and wait time have been reduced.

You issue this command on the first cell:

```
cellcli -e list metriccurrent attributes name, metricValue, metricType
where objectType = IORM_CATEGORY and and merticObjectName = OLTP and name
like 'CT_IO.*'
```

You examine the output from the first cell which contains:

CT_IO_BY_SEC	0 MB/sec	Instantaneous
CT_IO_LOAD	1	Instantaneous
CT_IO_RQ_LG	1,172 IO requests	Cumulative
CT_IO_RQ_LG_SEC	0.0 IO/sec	Rate
CT_IO_RQ_SM	360,325 IO requests	Cumulative
CT_IO_RQ_SM_SEC	22.7 IO/sec	Rate
CT_IO_UTIL_LG	0%	Instantaneous
CT_IO_UTIL_SM	0%	Instantaneous
CT_IO_WT_LG	0 ms	Cumulative
CT_IO_WT_LG_RQ	0.0 ms/ request	Rate
CT_IO_WT_SM	0 ms	Cumulative
CT_IO_WT_SM_RQ	0.0 ms/ request	Rate

Which two sets of metrics would you use to determine whether the I/O performance has improved for the OLTP category for the duration of the one-hour measurement period?

- A. CT_IO_RQ_SM, CT_IO_RQ_LG, CT_IO_RQ_SM_SEC, and CT_IO_RQ_LG_SEC
- B. CT_IO_UTIL_SM and CT_IO_UTIL_LG
- C. CT_IO_RQ_SM_SEC and CT_IO_RQ_LG_SEC
- D. CT_IO_WT_SM, CT_IO_WT_LG, CT_IO_WT_SM_RQ, and CT_IO_WT_LG_RQ

Answer: D

NEW QUESTION 5

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

- A. Database Writer I/O always prioritized over all user I/O.
- B. 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 P and Q are still mapped to the same categories.
- C. 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 the interdatabase plans are unchanged.
- D. 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 interdatabase plans are unchanged, and consumer groups A and B are still mapped to the same category.
- E. Log Writer I/O and Control File I/O are always prioritized over all user I/O.

Answer: BE

Explanation: B: Rules in an interdatabase resource plan specify allocations to databases, not consumer groups.
 E: Redo and control file writes always take precedence. References:

NEW QUESTION 6

Which two completely prevent a Smart Scan from occurring?

- A. querying a table containing many chained rows
- B. querying a table containing many migrated rows
- C. performing a minimum or maximum function on an indexed column
- D. performing a Fast Full Index scan on a reverse key index
- E. referencing more than 255 columns from an OLTP compressed table in a query
- F. querying a table containing a JSON column

Answer: BE

Explanation: B: Migrated Rows is a special case of chained rows. Migrated rows still affect performance, as they do in conventional storage situations, but with the additional overhead of reducing the beneficial effects of "Smart Scan" in addition to increasing the number of I/Os
 E: Smart Scans - broadly speaking and ignoring edge cases - can only transport a maximum of 254 columns from a single (non-HCC) segment. Requesting more columns will simply disable Smart Scans for that segment.
 An interesting limitation to Exadata Smart Scans - if more than 254 columns from a table (not HCC compressed, more on that in moment) need to be projected, Smart Scans for that particular segment will be disabled and Exadata will fall back to conventional I/O. This means that the number of columns in the projection clause can make a significant difference to performance, since only Smart Scans allow taking advantage of offloading and particularly avoiding I/O via Storage Indexes.

NEW QUESTION 7

Which three statements are true concerning InfiniBand port and subnet monitoring on an X5 Database Machine?

- A. The Infiniband subnet master location can be determined by using the getmaster command run on any database server.
- B. The Infiniband subnet master location can be determined by using the getmaster command run on an IB switch.
- C. The InfiniBand port status may be displayed on the storage servers by using the LIST IBPORT command in the CELLCLI utility.
- D. The InfiniBand port status may be displayed on the Infiniband switches by using the LIST IBPORT command in the DBMCLI utility.
- E. The InfiniBand port status may be displayed on the database servers by using the LIST IBPORT command in the DBMCLI utility.
- F. The InfiniBand port monitoring is automatic on the database servers and is managed by Enterprise Manager.

Answer: BCE

Explanation: B: From any InfiniBand switch in the network (leaf switch or spine switch), log in as root and run the getmaster command to obtain the location of the master SM as follows:

getmaster

This command displays the host name or IP address and the IP address of the switch where the master SM is running.

C: CellCLI> list ibport - Will display InfiniBand configuration details

E: The DBMCLI utility is the command-line administration tool for configuring database servers, and managing objects in the server environment.

The LIST IBPORT command displays attributes for InfiniBand ports determined by the specified attributes and filters.

References: http://docs.oracle.com/cd/E80920_01/DBMMN/exadata-dbmcli.htm

https://docs.oracle.com/cd/E18476_01/doc.220/e18478/GUID-9FF8B5B0-3481-4B73-89D3-108CBD7EB989.htm

NEW QUESTION 8

Which two activities are supported on the storage servers in an Exadata X6 Database Machine?

- A. Upgrading the Storage Server software package using RPM
- B. Configuring secure shell user equivalency for the cellmonitor user
- C. Installing an alternative package manager
- D. Changing the root password
- E. Upgrading a device driver for hard disks when inserting a replacement disk after a hard disk failure

Answer: AC

NEW QUESTION 9

Which two must be true for a Smart Scan to occur on a table?

- A. Sessions querying the table must set cell_offload_processing = true.
- B. It must be stored in an ASM diskgroup with a 4MByte AU size.
- C. It must be heap organized.
- D. It must be accessed using direct path reads.
- E. The table must not be compressed.

Answer: AD

NEW QUESTION 10

Your X6 Exadata Database Machine is running Oracle Database 12c, and 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 wish to compress these tables without causing unacceptable performance overheads to the OLTP workload.

Which three are true regarding this requirement?

- A. Compression is performed on database servers when using row store compress advanced in an Exadata environment.
- B. Using row store compress advanced will compress the data more than if using column store compress for archive low.
- C. The compression method column store compress for archive high is the worst fit for this requirement.
- D. Compression is performed on Exadata Storage Servers when using row store compress advanced in an Exadata environment.
- E. Using row store compress advanced will compress the data less than if using column store compress for query low.

Answer: ABD

Explanation: A: Creating a Table with Advanced Row Compression

The following example enables advanced row compression on the table orders: CREATE TABLE orders ... ROW STORE COMPRESS ADVANCED;

B: ARCHIVE LOW compression (Exadata only), recommended for Archival Data with Load Time as a critical factor

NEW QUESTION 10

Which two statements are true about Exadata X5 cell metrics and alerts?

- A. Cell alerts on each storage server are accumulated in memory by the CELLSRV process and stored on a filesystem-based repository.
- B. Cell metrics are written to disk every hour by default.
- C. Cell metrics on each storage server are accumulated in memory by the CELLSRV process and written to a filesystem-based repository.
- D. Cell alerts on each server are accumulated in memory by the MS process and stored on a filesystem-based repository.
- E. Cell alerts are written to disk every hour by default.

Answer: BD

Explanation: Metrics are a series of measurements that are computed and retained in memory for an interval of time, and stored on a disk for a more permanent history.

On the storage servers, the CELLSRV process provides the majority of Oracle Exadata storage services and is the primary storage software component. One of its functions is to process, collect, and store metrics. The Management Server (MS) process receives the metrics data from CELLSRV, keeps a subset of metrics in memory, and writes to an internal disk-based repository hourly.

References: <http://www.oracle.com/technetwork/articles/servers-storage-admin/monitor-exadata-em12-2291964.html>

NEW QUESTION 13

An important application has been migrated to a database on an X5 Database Machine.

You are examining high-load SQL statements from this application, in an attempt to determine which ones will benefit from the Exadata Smart Scan capability.

Which three access paths always generate "cell single block physical read" or "cell multiblock physical read" requests?

- A. Index fast full scans executed in parallel
- B. Full table scans on heap organized tables executed in parallel
- C. Full table scans on heap organized tables executed serially
- D. Full table scans on index organized tables executed in parallel
- E. Index unique scan access by primary key to heap organized tables executed in parallel
- F. Index unique scan access by primary key to heap organized tables executed serially

Answer: ACF

NEW QUESTION 14

Which two statements are true about the Integrated Lights Out Manager (ILOM) on an Exadata x5 or x6 Database Machine?

- A. ILOM on database servers can use IPMI to send sensor alerts to an Enterprise Manager Agent.
- B. ILOM on a storage server uses SNMP to send sensor alerts to the cellsrv process on the same storage server.
- C. ILOM on a storage server uses IPMI to send sensor alerts to the MS process on the same storage server.
- D. ILOM on a storage server uses SNMP to send sensor alerts to the MS process on the same storage server.
- E. ILOM on database servers can use SNMP to send sensor alerts to an Enterprise Manager Agent.

Answer: CE

Explanation: C: ILOM supports the Intelligent Platform Management Interface (IPMI), which enables you to monitor and control your server platform, as well as to retrieve information about your server platform.

ILOM supports alerts in the form of IPMI Platform Event Trap (PET) alerts. Alerts provide advance warning of possible system failures.

E: Oracle ILOM supports the Simple Network Management Protocol (SNMP), which is used to exchange data about network activity.

SNMP functionality requires the following two components: References:<https://docs.oracle.com/cd/E19860-01/E21452/E21452.pdf>

NEW QUESTION 18

Batch and DSS workloads on your Exadata X6 Database Machine are causing performance problems for OLTP workloads at certain times of the day. There are five RAC databases performing OLTP I/O. Two of them also perform batch I/O. There is another RAC database that performs only DSS I/O. You are asked to resolve this problem so that the OLTP workload will not suffer when competing with the batch or DSS workloads. You decide to use the I/O resource manager. Which is the best way to solve this problem?

- A. Create a category plan with EXADCLI calling EXACLI to give the OLTP category the highest priority on all the cells, and use the Database Resource Manager to create the OLTP category and map all OLTP- oriented consumer groups in all databases to it.
- B. Use the Database Resource Manager to give the OLTP category the highest priority and use EXADCLI calling EXACLI to create the OLTP category on all the cells.
- C. Create a category plan with EXADCLI calling EXACLI to give the OLTP category the highest priority on all the cells, and use the Database Resource Manager to give the OLTP consumer group the highest priority on all databases.
- D. Create a database plan using EXADCLI calling EXACLI to give OLTP I/O the highest priority from all six databases.

Answer: C

Explanation: You can manage I/O resources based on categories by creating a category plan.

You can add any number of categories, or modify the predefined categories. You should map consumer groups to the appropriate category for all databases that use the same cell storage. Any consumer group without an explicitly specified category defaults to the OTHER category.

References:http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server- iorm.htm

NEW QUESTION 23

Identify three valid reasons for creating multiple griddisks on a single harddisk-based celldisk.

- A. to implement storage realms so that storage that can be reserved for specific resource consumer groups in the same database
- B. to enable the creation of normal or high redundancy ASM diskgroups
- C. to segregate storage into multiple pools with different performance characteristics
- D. to enable disk mirroring for the system area
- E. to implement storage realms so that storage can be reserved for specific databases
- F. to implement storage realms so that storage that can be reserved for specific Grid Infrastructure

Answer: BCF

Explanation: Creating multiple grid disks per cell disk allows you to create multiple pools of storage on the same Exadata Storage Server. The multiple grid disks can be assigned to separate ASM diskgroups, which can be provisioned to different databases.

Note: Griddisk is a logical disk that can be created on a celldisk. In a standard Exadata deployment we create griddisks on hard disk based celldisks only. While it is possible to create griddisks on flashdisks, this is not a standard practice.

F: After you complete the cell configuration, you can perform the following optional steps on the storage cell:

References:https://docs.oracle.com/cd/E80920_01/SAGUG/SAGUG.pdf

NEW QUESTION 27

Which three statements are true regarding the configuration of Auto Service Request (ASR) on an X5 Database Machine?

- A. ASR Manager may be installed on any type of server running Oracle Solaris.
- B. ASR Manager may be installed on any type server running Oracle Linux.
- C. Oracle Linux must be used on the server where ASR Manager is installed.
- D. SMTP must be enabled on at least one database server.
- E. Oracle Solaris must be used on the server where ASR Manager is installed.
- F. ASR Manager must be installed on one of the database servers.
- G. HTTPS connectivity must be enabled from the ASR Manager host to the Internet.

Answer: ABG

Explanation: AB: Oracle recommends that you install Oracle ASR Manager on an external, standalone server. This server receives fault telemetry information from Oracle Exadata Database Machine servers. This server must run an Oracle Solaris or Linux operating system.

G: Before installing Oracle Auto Service Request (Oracle ASR) you must ensure connectivity to the Internet using HTTPS.

References: Oracle Auto Service Request Exadata Database Machine Quick Installation Guide

http://docs.oracle.com/cd/E80920_01/ASXQI/toc.htm#GUID-89192315-1B1A-4B5A-89BB-94BB5B6ABB05

NEW QUESTION 32

Which two are true about Smart Scan?

- A. A query rewrite may occur to a materialized view stored in Exadata, and it will always benefit from Smart Scan.
- B. Column projection does not contribute to the performance benefit of Smart Scan.
- C. A query rewrite may occur to a materialized view stored in Exadata but will never benefit from Smart Scan.
- D. It is possible to offload single row functions to the Exadata Storage Servers.
- E. Join processing can be partly offloaded to Exadata Storage Servers.
- F. Join processing can be fully offloaded to Exadata Storage Servers.

Answer: DE

Explanation: D: 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.

E: With Smart Scan processing, row filtering, column filtering, some join processing, and other functions are performed in the Exadata cells.

References: <http://www.centroid.com/blog/exadata-smart-scan-processing>

NEW QUESTION 33

You plan to partition the database and storage grids in an X5-2 full rack, creating two clusters and two storage grids without using virtualization.

One cluster will be used for production and should consist of six database servers and nine storage servers.

The other cluster will be used for test and development, and should consist of two database servers and five storage servers.

The storage must be partitioned so that the storage servers are visible only to the appropriate database servers that are meant to access them.

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 CELLI
- D. ORA file on each database server to contain IP addresses of storage servers in the storage grid associated with the cluster to which that database server belongs.
- E. Edit the CELLINI
- F. ORA file on each database server to contain IP addresses of storage servers in the storage grid associated with the cluster to which that database server belongs.
- G. Edit the CELLI
- H. ORA file on each database server to contain IP addresses of database servers which are allowed access to specific storage servers in the same storage grid.
- I. Edit the CELLI
- J. ORA file on each storage server on contain IP addresses of database servers in the database server grid that are associated with the storage grid to which that storage server belongs.

Answer: C

Explanation: The cellinit.ora file is host-specific, and contains all database IP addresses that connect to the storage network used by Oracle Exadata Storage Servers. This file must exist for each database that connect to Oracle Exadata Storage Servers.

The cellinit.ora file contains the database IP addresses. The cellip.ora file contains the storage cell IP addresses. Both files are located on the database server host.

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-configuring.htm#SAGUG20369

NEW QUESTION 37

You must replace the battery on the disk controller in three Exadata storage servers and must power off the affected storage servers.

Which two commands must you execute to safely power off the storage servers in an X5 Database Machine?

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

Answer: CE

Explanation: C: 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.

E: 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.

NEW QUESTION 42

You are planning your deployment of Enterprise Manager to monitor all the components of an X5 Database Machine. A part of the requirement is to provide for high availability of the monitoring infrastructure. If the host running the agent that has Database Machine targets bound to it fails, the monitoring of these targets must be done by another agent. Which three statements are true regarding the configuration used to support this requirement?

- A. Database Machine plug-ins must be deployed to at least two Enterprise Manager agents.
- B. Fail back to the original agent when the host is restarted is done automatically.
- C. Fail over to any secondary agent is done automatically.
- D. Fail over to any secondary agent must be done manually.
- E. Database machine plug-ins must be deployed to all Enterprise Manager agents.
- F. Fail back to the original agent when the host is restarted must be done manually.

Answer: BCE

NEW QUESTION 45

You plan to migrate a database supporting both DSS and OLTP workloads to your new X5 Database machine. The workloads contain many complex aggregating functions and expensive joins on large partitioned tables in the DSS workload and indexed access for OLTP workloads. Which three benefits accrue as a result of this migration?

- A. Superior compression capability designed specifically for OLTP workloads
- B. Columnar storage capability for data in row major data blocks that is held in flash cache
- C. Superior compression capability designed specifically for data warehouse tables
- D. Superior compression capability designed specifically for archival data
- E. Superior flash cache compression technique
- F. Cell offload processing for indexed-organized table access

Answer: ACD

Explanation: There are 6 different kinds of Table Data Compression methods:

NEW QUESTION 48

You plan to migrate an Oracle database that supports an online transaction processing (OLTP) workload to your X6 Database Machine. The Database Machine database version is 11.2. You plan to perform a physical database migration using Transportable Database. Which two are requirements for this method?

- A. The source database must be deployed on a Little Endian platform,
- B. The source database must be at least 11.1.
- C. The source database must be at least 11.2.
- D. The source database must use only bigfile tablespaces.
- E. The source database must be deployed on a Big Endian platform.
- F. The source database must not use bigfile tablespaces.

Answer: AC

Explanation: Transportable database approach (TDB). If the source system is running Oracle 11.2 or above with Little Endian format, this method can be used to migrate the database to Exadata.

References:<https://www.toadworld.com/platforms/oracle/w/wiki/11551.managing-troubleshooting-exadata-part-3-migrating-databases-to-exadata-database-machine-best-practices>

NEW QUESTION 50

You plan to migrate an existing production database supporting online transaction processing (OLTP) workloads to an X6 Exadata Database Machine. The database currently supports an application requiring fast response times to satisfy stringent business requirements, and most of the application queries use indexed access to tables. For which two cases would you consider dropping indexes that are not used for constraints after the migration to assure that Smart Scans occur?

- A. if Smart Scan performs better than any type of index scan on the corresponding table.
- B. if Smart only occur instead of index skip scans on the corresponding table.
- C. if Smart only occur instead of index range scans on the corresponding table.
- D. if Smart Scans performs equally well to any type of index scan on the corresponding table.

Answer: AC

NEW QUESTION 54

You plan to consolidate your company's INVENTORY and SALES databases onto your new Exadata X6 Database Machine. You are considering consolidation of all schemas into one RAC database. Which three factors would you need to consider before choosing this approach?

- A. the number of database users
- B. Security Policies
- C. the number of database roles
- D. the number of simultaneous sessions
- E. the maintenance schedule

Answer: ABD

NEW QUESTION 56

You are examining your Exadata storage servers for routine maintenance requirements and run the imageinfo command as shown:

```
[root@excel04 ~] # imageinfo
Kernel version: 2.6.18-194.3.1.0.4.el5 #1 SMP Sat Feb 19 03:38:37 EST 2011 x86_64 Cell version: OSS_11.2.0.3.0_LINUX.X64_110429.1
Cell rpm version: cell-11.2.2.3.1_LINUX.X64_110429.1-1
Active image version: 11.2.2.3.1.110429.1
Active image activated: 2011-05-09 16:26:36-0400 Active image status: success
Active image partition on device: /dev/md6 Active software partition on device: /dev/md8
In partition rollback: Impossible
Cell boot usb partition: /dev/sdm1
Cell boot usb version: 11.2.2.3.1.110429.1
Inactive image version: 11.2.1.2.6
Inactive image activated: 2010-10-04 23:59:16 -0400 Inactive image STATUS: success
Inactive system partition on device: /dev/md5
Inactive software partition on device: /dev/md7
Boot area has rollback archive for the version: 11.2.1.2.6 Rollback to the inactive partitions: Possible
Which two statements are true about the software and system partitions?
```

- A. /dev/md6 and /dev/md8 have one of their mirror copies on physical disk 1
- B. /dev/md6 and /dev/md8 are located only on physicaldisk1
- C. /dev/md5 and /dev/md7 have one of their mirror copies on physicaldisk 0
- D. /dev/md5 and /dev/md7 are located only on physicaldisk 0

Answer: BC

NEW QUESTION 57

Which two are benefits of an active/inactive configured InfiniBand network on Exadata Database Machine X5?

- A. Improved performance for Oracle Network traffic
- B. Improved reliability for Cache Fusion RAC network traffic
- C. Improved reliability when executing Distributed Command Line Interface (DCLI) to run CELLCLI commands
- D. Improved performance for ASM rebalance network traffic
- E. Improved performance when executing Distributed Command Line Interface (DCLI) to run CELLCLI commands

Answer: BC

Explanation: Active-passive bonding provides reliability through failover.

NEW QUESTION 60

Which two are true concerning Columnar Flash Caching?

- A. It can be enabled or disabled for individual databases by using I/O Resource Manager database plans.
- B. It automatically transforms ROW STORE COMPRESS compressed data into a purecolumnar format during Exadata Smart Flash Cache population.
- C. Data is either cached in Hybrid Columnar Compressed format or in pure compressed format, but never in both formats at the same time.
- D. It is enabled by default.
- E. It automatically transforms Hybrid Columnar Compressed (HCC) data into a pure columnar format during Exadata Smart Flash Cache population.
- F. It improves single-row lookup performance.

Answer: DE

Explanation: D: In-Memory Columnar Caching on cells is enabled by default when the INMEMORY_SIZE is configured. You do not need to do anything to get this enhancement.

E: Columnar Flash Caching implements a dual format architecture in Exadata flash by automatically transforming frequently scanned Hybrid Columnar Compressed data into a pure columnar format as it is loaded into the flash cache.

NEW QUESTION 61

Which two statements are true about data paths used during RMAN backup and restore operations when media management servers use InfiniBand to connect to an X5 Database Machine?

- A. During backups, data blocks are always read from Smart Flash Cache by cellsvr.
- B. During restores, data blocks read from the media are sent to a database server that sends the blocks to cellsvr to be written.
- C. During backups, data blocks are read by cellsvr and sent to a database server, which sends the blocks to the media manager to be written to media.
- D. During backups, data blocks are read by cellsvr and sent directly to the media manager to be written to media.
- E. During restores, data blocks read from the media are sent directly to cellsvr to be written.

Answer: DE

Explanation: The fastest database backup is achieved via InfiniBand.

RMAN does not back up directly to tape. However, it will integrate with media management software such as Oracle Secure Backup and utilize their capabilities to manage tape libraries.

Exadata Database Machine: The Database Machine contains the databases that need to be backed up. Oracle RMAN is the only mechanism to back up the databases that utilize Exadata Storage Servers as the storage. RMAN processes run on the database servers and interact with the Oracle Secure Backup (OSB) agent, which further interacts with the media management software and enables RMAN to communicate with the tape library.

References:http://apprize.info/data/oracle_4/6.html

NEW QUESTION 62

Which four statements are true about the configuration of Auto Service Request (ASR) for use with an X5 Database Machine?

- A. ASR Manager software may be installed on one of the database servers in the Database Machine.
- B. ASR can connect to the Internet using Simple Network Management Protocol (SNMP), using a proxy server.
- C. ASR Manager can connect to the Internet directly with HTTPS.
- D. ASR Manager can connect to the Internet with HTTPS using a proxy server.
- E. ASR can connect to the Internet directly using Simple Network Management Protocol (SNMP).
- F. ASR Manager software may be installed on a server that is not part of the Database Machine.

Answer: ACDF

Explanation: C: If you need to use HTTPS for security purposes, you can set up HTTPS/SSL for the ASR Manager HTTP receiver.

D: As part of the registration process for ASR Manager, you can optionally set the ASR Manager to access the internet through a proxy server.

References: https://docs.oracle.com/cd/E37710_01/install.41/e18475/ch2_asr_manager.htm#ASRUD137

NEW QUESTION 65

Which three statements are true about the CELLCLI command?

- A. It can execute commands on multiple storage servers in parallel.
- B. It has command-line history.
- C. It requires root privileges to create CELLDISKS and GRIDDISKS.
- D. It can be executed on storage servers.
- E. It can be executed using the DCLI utility.

Answer: BDE

Explanation: The CellCLI utility supports command-line history.

Each Exadata Storage cell can be monitored, configured, and maintained using the cellcli command line interface.

The dcli utility facilitates centralized management across an Oracle Exadata Storage Server Software realm by automating the execution of CellCLI commands on a set of cells and returning the output to the centralized management location where the dcli utility was run.

References: http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-cellcli.htm

NEW QUESTION 70

Which three are true concerning Exadata Smart Flash Cache on X6 Database Machines with High Capacity Storage Servers?

- A. The "keep" cache is limited to a maximum of 80% of the total Smart Flash Cache capacity.
- B. Smart Scan will always use the Smart Flash Cache for all I/Os fast full index capacity.
- C. Smart Scan will always use the Smart Flash Cache for all I/Os for full table scans.
- D. Single block reads can benefit from Smart Flash Cache.
- E. Multiblock reads can benefit from Smart Flash Cache.

Answer: BDE

Explanation: D: By default Exadata stores only small I/Os in the Exadata Smart Flash Cache. Small I/Os in most cases are single-block reads. During a full table scan Oracle requests blocks in multiblock lumps (by default 16 blocks), and these are

NEW QUESTION 73

The InfiniBand switches on your X5 Database Machine contain the software version as confirmed using the version command:

```
[root@exampsw -ib2 bin] # version SUN DCS 36p version: 1.3.3-2 Build time: Apr 4 2015 11:15:19
```

SP board info:

```
Manufacturing Date: 2014.05.05 Serial Number: "NCD3X0178" Hardware Revision: 0x0006 Firmware Revision: 0x0102 BIOS version: NOW1R112 BIOS date: 04/24/2014
```

Which two tools are useful for monitoring the InfiniBand switches for fan problems and by which users should they be run?

- A. Run the env_test command on the switches after logging in as the ilom-admin user.
- B. Run the getfanspeed command on the switches after logging in as root.
- C. Run the env_test command on the switches after logging in as root.
- D. Run the showtemps command on the switches after logging in as the ilom-admin user.

Answer: CD

Explanation: C: Login as root on IB switch and run,

```
# env_test
```

Example output include:

```
Starting FAN test: Fan 0 not present
```

```
WARNING Fan 1 running slow rpm 5450 limit 6440
```

```
WARNING Fan 2 running slow rpm 5341 limit 6440
```

```
WARNING Fan 3 running slow rpm 5341 limit 6440
```

D: The showtemps command displays internal temperatures for the switch.

The following example shows how to display switch temperatures with the showtemps command.

```
# showtemps
```

```
Back temperature 29
```

```
Front temperature 30
```

SP temperature 36
Switch temperature 52, maxtemperature 56 All temperatures OK

References:<http://myoracleidiary.blogspot.se/2014/04/all-about-infiniband-switches-on-exadata.html>

NEW QUESTION 74

Which two statements are true about Auto Service Request (ASR) with an X5 Database Machine?

- A. ASR Manager must be installed and configured on a dedicated server external to the Database Machine.
- B. Configuring ASR is mandatory for all Database Machine assets.
- C. ASR Manager opens a service request (SR) automatically after sensors detect hardware faults.
- D. ASR Manager must be installed and configured on one of the database servers.
- E. ASR can upload configuration metadata to support problem resolution.
- F. ASR communicates with Oracle support services using HTTPS.

Answer: AC

Explanation: A: Oracle recommends that you install Oracle ASR Manager on an external, standalone server.

C: Oracle ASR is designed to generate Oracle service requests automatically when certain types of faults are detected on Oracle products that are qualified for Oracle ASR.

Oracle ASR works only for specific component faults. Most of the common components, such as disks, fans, and power supplies, are covered. However, some components are not covered. For example, Oracle ASR does not cover InfiniBand events; there are specific images and specific InfiniBand switch firmware that you must use. You cannot upgrade these components independently.

NEW QUESTION 79

You issued these commands to all Exadata Storage Servers in an X6 Exadata Database Machine using dcli:

```
alter iormplan objective = low_latency alter iormplan active
```

There are no database or category plans defined.

You are encountering disk I/O performance problems at certain times, which vary by day and week.

DSS and Batch workloads perform well some of the time.

Further investigation shows that at times, the workloads are all OLTP I/Os, at other times all batch I/Os, and sometimes a bit of each.

You wish to have disk I/O managed so that performance will be optimized for all workloads. Which statements would you issue to all Exadata Storage Servers to achieve this?

- A. alter iormplan objective=high_throughput
- B. alter iormplan objective=balanced
- C. alter iormplan objective=low_latency
- D. alter iormplan objective= auto
- E. alter iormplan objective= ' '

Answer: D

Explanation: The supported IORM objectives are auto, low_latency, balanced, and high_throughput. The recommended objective option is auto which allows IORM to continuously monitor the workloads, and select the best mode based on the active workloads currently on the cells.

References:http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-iorm.htm

NEW QUESTION 80

You plan to monitor the status of the motherboard, memory, power, fans, and network cards on the database nodes in your Exadata X6 Database Machine using Enterprise Manager.

Where must you set the thresholds for these hardware components and why, to assure that sensor readings, faults, and any related alerts, are visible in Enterprise Manager?

- A. No thresholds need to be set because they are present in the ILOM and in Enterprise Manager.
- B. Set thresholds in ILOM and in Enterprise Manager because they are not present anywhere and must be set in both places.
- C. Set thresholds only in ILOM because they are not preset anywhere but need to be set only in ILOM.
- D. No thresholds need to be set because they are preset in the ILOM and these are sufficient for monitoring.

Answer: D

NEW QUESTION 84

You have partitioned an X5-2 full rack into two four-node RAC clusters called CLUSA and CLUSB.

No virtualization is used and the storage grid has not been partitioned.

Which two set of files on which servers must be modified after connecting an Exadata storage expansion rack to your X5-2 Exadata Database Machine on the InfiniBand network so that the storage servers in the expansion rack are added to the storage grid?

- A. the CELLINIT.ORA files on the 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 Exadata storage servers
- E. the CELLIP.ORA files on the database servers in CLUSB
- F. the CELLIP.ORA files on all newly added Exadata storage servers

Answer: BE

Explanation: The cellip.ora is the configuration file, on every compute node (database server), that tells ASM instances which cells are available to this cluster.

The cellip.ora file contains the storage cell IP addresses. References:http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-configuring.htm#SAGUG20369

NEW QUESTION 87

Identify two supported modifications to an X5 Database Machine.

- A. installing a Host Bus Adapter (HBA) in the compute nodes
- B. installing additional Operating System (O/S) packages on the computer nodes
- C. installing additional Operating System (O/S) packages on the Exadata storage servers
- D. installing a Host Bus Adapter (HBA) in the storage servers
- E. replacing the Cisco Ethernet switch with another make or model
- F. replacing the Infiniband switches with another make or model

Answer: BD

Explanation: B: Updating Exadata Database Servers (compute nodes).

The software and firmware components that are updated for a specific release depend on the current Exadata software release the database server is running and the release you are updating to. Linux operating system packages and Exadata software are always updated while firmware may be updated for only a small selection of the components or not at all.

D: The storage cells (storage servers/nodes) contain HBAs, and spare PCI I/O slots. References:https://docs.oracle.com/cd/E41033_01/html/E55032/bbgjapie.html
http://docs.oracle.com/cd/E80920_01/DBMMN/updating-exadata-software.htm#DBMMN-GUID-AB25ED9A-7920-441A-9A60-14ED2753B01C

NEW QUESTION 90

Which four are true about Exadata features?

- A. Storage Indexes persist across Exadata storage server reboots.
- B. Data Warehouse workloads will benefit from Smart Flash Cache configured in Write-Through mode.
- C. Hybrid Columnar Compressed tables can be compressed and decompressed on Exadata storage servers.
- D. Hybrid Columnar Compressed tables can be compressed and decompressed on the database servers.
- E. OLTP workloads will benefit from Smart Flash Cache configured in Write-Back mode.
- F. Storage Indexes persist across database server reboots.

Answer: CDEF

Explanation: CD: That data remains compressed not only on disk, but also remains compressed in the Exadata Smart Flash Cache, on Infiniband, in the database server buffer cache, as well as when doing back-ups or logshipping to Data Guard.

E (not B): Use the Write-Back Flash Cache feature to leverage the Exadata Flash hardware and make Exadata Database Machine a faster system for Oracle Database Deployments. Write-through cache mode is slower than write-back cache mode. However, write-back cache mode has a risk of data loss if the Exadata Storage Server loses power or fails.

F: Storage indexes are not stored on disk; they are resident in the memory of the storage cell servers.

NEW QUESTION 94

In which scenario will griddisks residing on the same celldisk have balanced performance characteristics in an X5 High Capacity Exadata Storage Server?

- A. griddisks on memory based celldisks
- B. griddisks on hard disk-based celldisk consumed by ASM when creating diskgroups using ASM-based Intelligent Data Placement (IDP)
- C. griddisks on hard disk-based celldisks consumed by the Database Smart Flash Cache
- D. griddisks on Flash-based celldisks

Answer: B

Explanation: Intelligent Data Placement enables you to specify disk regions on Oracle ASM disks for best performance. Using the disk region settings, you can ensure that frequently accessed data is placed on the outermost (hot) tracks which have greater speed and higher bandwidth. In addition, files with similar access patterns are located physically close, reducing latency. Intelligent Data Placement also enables the placement of primary and mirror extents into different hot or cold regions.

References:https://docs.oracle.com/cd/E11882_01/server.112/e18951/asmdiskgrps.htm#O_STMG10072

NEW QUESTION 97

Which three are true concerning Exadata snapshot databases?

- A. They are supported on non-container databases.
- B. They are based on a read-write copy of an existing database.
- C. They are integrated with the Multitenant architecture.
- D. They can be created only on sparse ASM disk groups.
- E. They don't support all Exadata features.
- F. They can be created on any type of ASM disk group.

Answer: ACD

Explanation: A: An Exadata snapshot database can be either a non-container database (non-CDB) or a container database (CDB). Creating an Exadata snapshot database of a CDB enables access to all of the pluggable databases in that container.

C: You can create two types of Exadata snapshots, depending on the current setup of your environment:

D: SPARSE disk group based database snapshots is functionality included in Exadata Storage Software. It requires Exadata Storage Software version 12.1.2.1.0 and Oracle Database 12c version 12.1.0.2 with bundle patch 5 or later. This feature is designed to work on native Exadata ASM storage disk groups. It uses ASM

SPARSE grid disk based thin provisioning where snapshot databases created on a SPARSE disk group need only the space for changes plus some metadata, thereby enabling storage efficient snapshot databases.

References:

http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-snapshots.htm#SAGUG-GUID-E1D6EF45-36EF-40E3-A57E-F80B749E6122

<http://www.oracle.com/technetwork/database/exadata/learnmore/exadata-database-copy-twp-2543083.pdf>

NEW QUESTION 99

An Exadata storage server physical disk on an X5 high-capacity full rack entered the predictive failure state. Identify the two steps that you must perform to replace this failed physical disk.

- A. Add the griddisks back into the ASM diskgroup they used to be a member of.
- B. Create a new celldisk and new griddisks on the replaced physical disk.
- C. Verify that the griddisks located on the physical disk have been successfully dropped from the associated ASM diskgroups.
- D. Identify the griddisks located on the failed physical disk and drop them from the associated ASM diskgroups.
- E. Replace the failed physical disk.

Answer: CE

Explanation: You may need to replace a physical disk because the disk is in warning - predictive failure status. The predictive failure status indicates that the physical disk will soon fail, and should be replaced at the earliest opportunity. The Oracle ASM disks associated with the grid disks on the physical drive are automatically dropped, and an Oracle ASM rebalance relocates the data from the predictively failed disk to other disks.

Note: After the physical disk is replaced, the grid disks and cell disks that existed on the previous disk in that slot are re-created on the new physical disk. If those grid disks were part of an Oracle ASM group, then they are added back to the disk group, and the data is rebalanced on them, based on the disk group redundancy and ASM_POWER_LIMIT parameter.

References: http://docs.oracle.com/cd/E80920_01/DBMMN/maintaining-exadata-storage-servers.htm#DBMMN21047

NEW QUESTION 100

Your Exadata X6-2 Database Machine has the exachk utility preinstalled. You decide to use it periodically to validate the installation against Oracle's recommended best practices.

Which three actions could you take to do this?

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

Answer: ADE

Explanation: AD: To schedule exachk, we can create a cron job or we can create a job in Enterprise Manager..

To add an Enterprise Manager job to execute Exachk on a periodic basis. Execute the following steps. First go to Enterprise, Job, Library. Then create a new OS Command job

E: exachk options include AUTORUN_SCHEDULE, which schedules recurring health checks using the AUTORUN_SCHEDULE daemon option.

References:

<http://www.oracle.com/technetwork/database/availability/exadata-health-resource-usage-2021227.pdf>

http://docs.oracle.com/cd/E68491_01/OEXUG/automated-daemon-mode-operation.htm

Best practices: Execute the latest version of the exachk utility before and after planned maintenance activities and at least once every month

NEW QUESTION 102

Which three factors should you consider when choosing a method for migrating a database to the X5 Database Machine?

- A. The down time allowed for the migration
- B. Endian format of the source database
- C. Number of tablespaces in the source database
- D. The type of database workloads
- E. Size of the source database

Answer: BDE

Explanation: B: Endian format: Non-Exadata format

D: Use Real Production Workload

– Real Application Testing (RAT) References:

NEW QUESTION 107

Which two statements are true about the Enterprise Manager plug-in configuration for the X5 Database Machine?

- A. There is one plug-in for Cloud Control.
- B. There are several separate plug-ins for Cloud Control.
- C. Some components of the Database Machine require SNMP trap forwarders.
- D. All components of the Database Machine require SNMP trap forwarders.

Answer: AC

NEW QUESTION 109

Which three are true concerning Storage Indexes?

- A. A maximum of eight table columns for any table are indexed per storage region.
- B. The use of Storage Indexes for particular categories of I/O can be displayed by using an I/O Resource Manager Category Plan.
- C. Storage Indexes persist across any Exadata storage server reboots.
- D. The use of Storage Indexes for a particular database can be disabled by using an I/O Resource Manager Database Plan.
- E. A Storage Index is automatically maintained by CELLSRV based on the filter columns of the offloaded SQL.
- F. Different storage regions may have different columns indexed for the same table.

Answer: ADE

Explanation: A: Each disk in the Exadata storage cell is divided into equal sized pieces called storage regions (default 1MB). There is an index entry for every storage regions (1MB of data stored on disk). Each entry contains the minimum and maximum value for columns seen in 'where' clause predicates. Information for up to 8 columns can be stored. The index is then used to eliminate disk IO by identifying which storage regions don't match the 'where' clause of a query.

NEW QUESTION 112

You are patching your Exadata X6 Database Machine by applying a new image to the Storage Servers in a rolling fashion.

Your ASM environment on the Database Machine has five diskgroups stored on an unpartitioned Exadata storage grid, with redundancy settings as shown:

1. DATA_DBM1 – Normal Redundancy
2. RECO_DBM1 – Normal Redundancy
3. DBFS_DG – Normal Redundancy
4. DATA2_DBM1 – High Redundancy
5. DATA3_DBM1 • High Redundancy

Which two diskgroups will not suffer from any data loss throughout the patching process even if there is a single disk failure on one of the cells

- A. DBFS_DG
- B. DATA3_DBM1
- C. DATA2_DBM1
- D. DATA_DBM1
- E. RECO_DBM1

Answer: DE

Explanation: HIGH redundancy provides protection against 2 simultaneous disk failures from 2 distinct storage servers or 2 entire storage servers. HIGH redundancy provides redundancy during Exadata storage server rolling upgrades.

References:<http://blog.umairmansoob.com/choosing-high-vs-normal-redundancy-with-exadata/>

NEW QUESTION 116

Which three are true concerning Hybrid Columnar Compression (HCC) deployed on Exadata storage?

- A. HCC can be used only when the Exadata Smart Flash Cache is configured in Write- Through mode.
- B. HCC data is never cached in the Exadata Smart Flash Cache.
- C. By default, decompression is performed by Exadata Storage Servers.
- D. Rows residing in HCC compressed segments, are always self- contained in a single database block.
- E. Rows residing in HCC compressed segments, are always self- contained in a single compression unit.
- F. Row-level locks are supported on HCC compressed data.

Answer: CEF

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

E: A logical construct called the compression unit is used to store a set of hybrid columnar compressed rows. When data is loaded, column values for a set of rows are grouped together and compressed. After the column data for a set of rows has been compressed, it is stored in a compression unit.

F: What happens when I update a row on compressed tables? What about locks?

Note: Oracle's Hybrid Columnar Compression technology is a new method for organizing data within a database block. As the name implies, this technology utilizes a combination of both row and columnar methods for storing data. This hybrid approach achieves the compression benefits of columnar storage, while avoiding the performance shortfalls of a pure columnar format.

NEW QUESTION 117

You are evaluating the performance of a SQL statement that accesses a very large table. You run this query:

```
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 10%');
```

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

Identify two reasons why the “cell physical IO interconnect bytes” statistic is greater than the “cell physical IO interconnect bytes returned by smart scan” statistic.

- A. There is a transaction, which committed after the query began, 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.
- B. There are chained rows in the table, causing some “single block physical reads” to be requested by the database instance, resulting in additional I/O.
- C. The table is a hash clustered table, causing “cell multiblock physical reads” to be requested by the database instance, resulting in additional I/O.
- D. The table is list partitioned, causing “cell list of blocks physical reads” to be requested by the database instance, resulting in additional I/O.
- E. There is a local index on a list partitioned table on the column used in the WHERE clause, causing “cell list of blocks physical reads” to be requested by the database instance, resulting in additional I/O.

Answer: CD

Explanation: C: Scan on a clustered table can prevent a Smart Scan from occur.

D: Scan on an index-organized table can prevent a Smart Scan from occur.

Note: The Cell physical IO interconnect bytes returned by smart scan metric shows how many bytes of I/O were returned by a smart scan to the database server.

References:<https://uhesse.com/2011/01/19/exadata-part-i-smart-scan/>

NEW QUESTION 120

Which three statements are true about bulk data loading capabilities in an X5 Database Machine environment?

- A. DBFS must use the DBFS_DG diskgroup for any DBFS store.
- B. DBFS may be used if Exadata-based shared storage is required to stage data before bulk loading into a database.
- C. DBFS must be used to have a POSIX-compliant shared storage solution that is accessible from the database servers on a Database Machine.
- D. DBFS must be used to bulk load data into a production database on the Database Machine.
- E. ACFS may be used if Exadata-based shared storage is required to stage data before bulk loading into a database.
- F. ACFS must be used to have a POSIX –compliant shared storage solution that is accessible from the database servers on a Database Machine.
- G. ACFS may use the DBFS_DG diskgroup to contain the ADVN volume file.

Answer: BCD

Explanation: External tables on DBFS file-systems provide the probably the most high-performance way to bulk load data into your database.

Bulk Data Loading

Describe the various options for staging data to be bulk loading into Database Machine Configure the Database File System (DBFS) feature for staging input data files

Use external tables based on input data files stored in DBFS to perform high-performance data loads

NEW QUESTION 124

You must apply patches and patch bundles in a rolling fashion, if possible, on the components of your Exadata X6 Database Machine.

You use RAC for your databases and also use Data Guard, having standby databases on another Database Machine.

You want to have scripts that contain the appropriate commands to patch your environment.

Your patch bundle is one the first database server and is located at /u01/stage. You have downloaded the oplan utility to the first database server and run:

```
$ORACLE_HOME/oplan/oplan generateApplySteps /u01/stage
```

Which two statements are true concerning the type of instructions oplan will generate?

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

Answer: AC

Explanation: Oplan generates instructions for all of the nodes in the cluster. Note:

NEW QUESTION 125

You issued these commands to all Exadata Storage Servers in an X6 Exadata Database Machine using dcli:

alter iormplan objective = off alter iormplan active
 There are no database or category plans defined.

You are suffering I/O performance problems at certain times, which vary by day and week. DSS and batch workloads perform well, but your OLTP workloads suffer poor response

times when running at the same time as these other workloads.

You must ensure that DSS and batch workloads retain relatively high throughput without causing excessive degradation of OLTP performance.

Which statement would you issue to all Exadata Storage Servers to achieve this?

- A. alter iormplan objective = low_latency
- B. alter iormplan objective = high_throughput
- C. alter iormplan objective = balanced
- D. alter iormplan objective = auto

Answer: A

Explanation: When the objective option is set to basic, the database resource plan maximum utilization limits are not enforced. For stricter plan conformance, and enforcement of maximum utilization limits, the objective option must be set to something other than basic. The supported IORM objectives are auto, low_latency, balanced, and high_throughput.

The objective may also be changed to low_latency in which case the latency for critical I/Os is very good but there is a significant degradation in scan throughput when both workloads are running concurrently.

NEW QUESTION 129

Which two statements are true about backup performance when using Recovery Manager (RMAN) on an X5 Database Machine?

- A. Backups are fast because the Infiniband network is faster than fiber channel storage, so RMAN read I/O operations are faster than in traditional storage networks.
- B. Backups are fast, because physical I/Os operations are reduced by the use of the Storage Index by cellsrv.
- C. Backups are fast, because data blocks are read from the Database Flash Cache.
- D. Backups are fast, because all data blocks are read from the Exadata Smart Flash Cache.
- E. For incremental backups using Block Change Tracking, cellsrv filters blocks and returns only those that have changed since the last backup.

Answer: DE

Explanation: D: The storage server software manages the flash cache and can recognize different types of I/O requests so that non-repeatable data access like RMAN backup I/O does not flush database blocks from the cache. It also prioritizes frequently accessed block types such as redo logs, control files, and index root blocks.

E: RMAN block change tracking allows incremental backups to run very quickly and efficiently. With block change tracking, only the areas of the database that have been modified since the last incremental backup or full backup are read from disk.

NEW QUESTION 134

Which two are true about sparse griddisks and their use in disk groups on an X5 Exadata Database Machine?

- A. Sparse diskgroups must be created using sparse griddisks.
- B. Sparse diskgroups may be created using a combination of sparse and non-sparse griddisks.
- C. Sparse diskgroups may not be used for database snapshots.
- D. Additional space for a sparse griddisk is allocated as soon as newly written data is stored in the flashcache on a cell.
- E. The virtual size of a sparse griddisk may exceed the physical size of the space occupied by the griddisk.

Answer: AE

Explanation: A: A sparse ASM disk group is composed of sparse grid disks.

E: Sparse grid disks allocate space as new data is written to the disk, and therefore have a virtual size that can be much larger than the actual physical size.

Sparse grid disks can be used to create a sparse disk group to store database files that will use a small portion of their allocated space. Sparse disk groups are especially useful for quickly and efficiently creating database snapshots on Oracle Exadata. Traditional databases can also be created using a sparse disk group.

References:

http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server- snapshots.htm#SAGUG-GUID-42945059-13FD-4F6A-B7FA-A1201D16238F

http://docs.oracle.com/cd/E80920_01/DBMSO/exadata-whatsnew.htm#DBMSO22120

NEW QUESTION 136

Which statement is true about operating systems in an X5 Database Machine multitrack configuration consisting of two full racks and one Exadata storage expansion rack?

- A. All Exadata storage servers used by the same virtual cluster nodes must run the same O/S but Exadata Storage Servers in different clusters may run different operating systems.
- B. All Exadata storage servers must run the Oracle Solaris O/S and all database servers within the same cluster must run Oracle Linux.
- C. All Exadata storage servers may run Oracle Virtual Machine (OVM).
- D. All Exadata storage servers must run Oracle Linux.
- E. All Exadata storage servers must run the Oracle Linux O/S and all database servers within the same cluster must run the same version of Oracle Virtual Machine (OVM).

Answer: D

Explanation: On both physical and virtual deployments, Exadata systems use minimal Linux distributions to ensure that just the RPMs needed to run Oracle database, are installed and enabled.

References: <http://www.oracle.com/technetwork/database/exadata/exadata-x5-2-ds- 2406241.pdf>

NEW QUESTION 140

You are designing the monitoring architecture and procedures for an X5 Database Machine fabric consisting of two full racks, to provide for the monitoring for all components on the critical path of the admin network.
 Which two components need to be monitored?

- A. the Infiniband leaf switches
- B. the power distribution units
- C. the Cisco 48 port Catalyst Ethernet Switch
- D. the InfiniBand spine switches

Answer: CD

Explanation: The Cisco Catalyst 4948 Ethernet switch is supplied with the Oracle Exadata Rack. D: Figure: Running the Subnet Manager in Different Rack Configurations

Rack Configuration	SM Should Run On...	SM Priority
Single Exalogic machine	All leaf switches	All leaf switches: 5
Two half- or full-rack Exalogic machines	Spine switches	Spine switch: 8
Two quarter-rack Exalogic machines	All leaf switches	All leaf switches: 5
Three or more Exalogic machines	Spine switches	Spine switch: 8
Half- or full-rack Exalogic machine connected to a half- or full-rack Exadata machine.	Spine switches	Spine switch: 8

References:

- https://docs.oracle.com/cd/E18476_01/doc.220/e18478/GUID-9FF8B5B0-3481-4B73-89D3-108CBD7EB989.htm#ELMOG76346
- http://docs.oracle.com/cd/E80920_01/DBMIN/configuring-exadata.htm#DBMIN21274

NEW QUESTION 143

You plan to migrate a database supporting an OLTP workload to your new X5 Database Machine.
 The current database instance supports a large number of short duration sessions and a very high volume of short transactions.
 Which three X5 Database Machine features can improve performance for this type of workload?

- A. An improved highly efficient undo and redo architecture
- B. Faster optimization due to an Exadata-specific optimizer
- C. Reduced I/O latency for writes due to writeback flashcache on all Exadata X5 and later models
- D. Ultra high I/O performance for reads and writes when using Exadata Extreme Flash in X5 and later models
- E. Reduced I/O latency for reads due to read flashcache on all Exadata X5 and later models

Answer: ACD

Explanation: A: To further accelerate OLTP workloads, the Exadata Smart Flash Cache also implements a special algorithm to reduce the latency of log write I/Os called Exadata Smart Flash Logging.

C: Use the Write-Back Flash Cache feature to leverage the Exadata Flash hardware and make Exadata Database Machine a faster system for Oracle Database Deployments.

D: Exadata X5-2 introduces Extreme Flash Storage Servers. Each Extreme Flash storage server contains eight 1.6 TB state-of-the-art PCI Flash drives. PCI flash delivers ultra-high performance by placing flash memory directly on the high speed PCI bus rather than behind slow disk controllers and directors.

References:<http://www.oracle.com/technetwork/database/exadata/exadata-x5-2-ds-2406241.pdf>

<http://www.oracle.com/technetwork/articles/database/exadata-write-back-flash-2179184.html>

NEW QUESTION 147

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 1z0-070 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 1z0-070 Product From:

<https://www.2passeasy.com/dumps/1z0-070/>

Money Back Guarantee

1z0-070 Practice Exam Features:

- * 1z0-070 Questions and Answers Updated Frequently
- * 1z0-070 Practice Questions Verified by Expert Senior Certified Staff
- * 1z0-070 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * 1z0-070 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year