

# Oracle

## Exam Questions 1z0-070

Oracle Exadata X5 Administration



#### 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](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 regarding the use of Auto Service Request (ASR) with an X6 Database Machine?

- A. The database server ILOMs must use SMTP over the management network for notifications to ASR Manager.
- B. The database server ILOMs must have SNMP traps configured to use the management network for notifications to ASR Manager.
- C. The storage server ILOMs must have SNMP traps configured to use the management network for notifications to ASR Manager.
- D. The database server MS process must have SNMP traps configured to use the management network for notifications to ASR Manager.

**Answer:** BC

**Explanation:** B: Database Server ILOM plug-in

Monitoring databases and their instances, ASM environments, the Grid Infrastructure, and the host software environment are done by Enterprise Manager in the usual way as these are standard targets. But monitoring the hardware for the database servers requires the ILOM plug-in, as there is no Management Server (MS) on the database servers to receive SNMP traps from the ILOM. The plug-in will receive sensor state and availability data from the ILOM including alerts based on pre-set ILOM thresholds.

C: Exadata Storage Server plug-in extends the monitoring of exadata cells in addition to providing a GUI interface. The plug-in uses an SSH connection to the cellmonitor user on the cells and uses list commands only. This is for interactive monitoring. One may also set thresholds using the plug-in which are distinct from any thresholds set using cellcli utility as the celladmin user. For alerts to be sent to the plug-in, SNMP traps are used as follows:

Cell ILOM alerts are sent to the cell Management Server (MS) via an SNMP trap. The MS then send SNMP notifications onward to the plug-in.

Cell alerts flagged by MS itself, such as cell thresholds being exceeded, or ADR software alerts, are sent to the plug-in using SNMP.

References: <https://dbatrain.wordpress.com/2011/06/>

[http://docs.oracle.com/cd/E21659\\_01/html/E21660/z40015671004046509.html](http://docs.oracle.com/cd/E21659_01/html/E21660/z40015671004046509.html)

#### NEW QUESTION 4

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 5

Which type or types of network traffic are transported over the internal InfiniBand interconnect in Exadata Database Machine X5?

- A. IDB protocol traffic, clustered ASM traffic, and clustered database instance traffic
- B. IDB protocol traffic and clustered database instance traffic only
- C. IDB protocol traffic and clustered ASM instance traffic only
- D. Both clustered ASM and clustered database instance traffic only
- E. IDB protocol traffic only

**Answer:** E

**Explanation:**

Oracle Exadata uses the Intelligent Database protocol (iDB) to transfer data between Database Node and Storage Cell Node. iDB is used to ship SQL operations down to the Exadata cells for execution and to return query result sets to the database kernel.

**NEW QUESTION 6**

Examine this IORM plan:

```
CellCLI> list iormplan detail

name: dmorlcel-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, limit=60,
         flashcache=on, flashlog=on
         name= finance, level=1, allocation=45, limit=60,
         flashcache=on, flashlof=off
         name=other, level=1, allocation=10,
         flashcache=off, flashlog=on

objective: balanced
status: active
```

Which two are true concerning this plan?

- A. The Finance database can use at least 45%, but never more than 60%, of the total Flash Cache capacity.
- B. I/Os from the finance database are guaranteed to get a minimum of 45% of the I/O bandwidth if the demand exists, and a maximum of 60% of the I/O bandwidth even if noother databases are doing I/O to the cell, and the demand from the finance database exceeds 60% of the maximum I/O rate of the cell.
- C. If I/Os come from the HR database only, then they may get up to 100% of the I/O bandwidth on the cell.
- D. I/Os from the finance database are guaranteed to get 45% of the I/O bandwidth if the demand is at least 60% of the maximum I/O rate of the cell, but may get 100% of the I/O bandwidth if no other databases are performing I/O to the cell.
- E. If I/Os come from the HR database only, then they may get up to 10% of the I/O bandwidth on the cell.

**Answer: BE**

**Explanation:** The IORM plan can be configured using the ALTER IORMPLAN command on command- line interface (CellCLI) utility on each Exadata storage cell. It consists of two parameters - dbplan and catplan. While the "dbplan" is used to create the I/O resource directives for the databases, the "catplan" is used to allocate resources by workload category consolidated on the target system.

allocation/share - Specify the resource allocation to a specific database in terms of percentage or shares.  
 limit - Specify maximum limit of disk utilization for a database.

**NEW QUESTION 7**

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 8**

You are checking the status of the ports on one of the InfiniBand switches on your X5 Database Machine.

You run the getportstatus command 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 the ibqueryerrors.pl script after logging in to the InfiniBand switch as root.

- B. Use Enterprise Manager Cloud Control.
- C. Use Enterprise Manager Express.
- D. Create user-defined metrics for the InfiniBand switch.

**Answer:** AB

**Explanation:** There are two approaches for monitoring Oracle Exadata Storage Servers: using a command-line interface (CLI) or using the graphical interface provided by the Oracle Enterprise Manager Cloud Control 12c console.

You can manually monitor the InfiniBand Switch Ports with the `ibqueryerrors.pl` command. Run the `ibqueryerrors.pl` command to report on switch port error counters and port configuration information using the command:

```
# ibqueryerrors.pl -rR -s RcvSwRelayErrors,XmtDiscards,XmtWait,VL15Dropped
```

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

#### NEW QUESTION 9

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

#### NEW QUESTION 10

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

#### NEW QUESTION 10

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 13

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 17

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 21

Which two statements are true about the use of Intelligent Platform Management Interface (IPMI) on an X5 or X6 Database Machine?

- A. IPMI can be used for server configuration and management on each database server.
- B. The Baseboard Management Controller (BMC), which supports IPMI, runs as a multithreaded process on the storage server O/S.
- C. The Baseboard Management Controller (BMC), which supports IPMI, runs as a single-threaded process on the storage server O/S.
- D. IPMI can be used for server configuration and management on each storage server.

**Answer:** AC

**Explanation:** A: IPMI – short for Intelligent Platform Management Interface – is an interface standard that allows remote management of a server from another using standardized interface. The servers in the Exadata Database Machine follow that. It's not an Exadata command but rather a general Linux one.

To power on a cell or database server, issue this from another server:

```
# ipmitool -H procel01-ilom -U root chassis power on
```

To stop a server, use the shutdown command. To stop immediately and keep it down, i.e. not reboot, execute: # shutdown -h -y now

C: Like the KCS interface, the SSIF Interface is only specified as a 'Single Threaded Interface' for standard IPMI commands. That is, the BMC implementation is not expected to process more than one IPMI request at a time. While an implementation is allowed to have a degree of 'command queuing', for standard IPMI messages the SSIF lacks a 'Seq' field that software can use to match up particular instances of requests with responses.

References: <https://www.intel.com/content/dam/www/public/us/en/documents/specification-updates/ipmi-intelligent-platform-mgt-interface-spec-2nd-gen-v2-0-spec-update.pdf>

#### NEW QUESTION 23

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 25

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 28

Which three statements are true concerning the configuration of SNMP on an Exadata x5 or x6 Database Machine?

- A. All SNMP notifications from the storage server ILOMs are sent directly to the Automatic Service Request (ASR) manager.
- B. Some SNMP notifications may be configured using the alter cell command on the storage servers.
- C. The Database Machine internal Cisco Switch communicates directly with the Enterprise Manager agent via SNMP.
- D. Some SNMP notifications from a storage server ILOM are sent directly to the management server (MS) process on the same storage servers.
- E. All SNMP notifications from a database server ILOM are sent directly to the Automatic Service Request (ASR) manager.
- F. All SNMP notifications from a database server ILOM are sent directly to the management server (MS) process on the same database server.

**Answer:** BDE

**Explanation:** B: You can configure SNMP Trap Destinations for the Exadata Storage Servers with the following command:

```
# cellcli -e "alter cell snmpSubscriber=(host ='ASR-Manager-name-or-IP- address',port=162,community=public,type=asr)"
```

D: Exadata Storage Server plug-in

This extends the monitoring of exadata cells in addition to providing a GUI interface. The plug-in uses an SSH connection to the cellmonitor user on the cells and uses list commands only. This is for interactive monitoring. One may also set thresholds using the plug-in which are distinct from any thresholds set using cellcli utility as the celladmin user. For alerts to be sent to the plug-in, SNMP traps are used as follows:

Cell ILOM alerts are sent to the cell Management Server (MS) via an SNMP trap. The MS then send SNMP notifications onward to the plug-in.

Cell alerts flagged by MS itself, such as cell thresholds being exceeded, or ADR software alerts, are sent to the plug-in using SNMP.

E: Database Server ILOM plug-in

Monitoring databases and their instances, ASM environments, the Grid Infrastructure, and the host software environment are done by Enterprise Manager in the usual way as these are standard targets. But monitoring the hardware for the database servers requires the ILOM plug-in, as there is no Management Server (MS) on the database servers to receive SNMP traps from the ILOM. The plug-in will receive sensor state and availability data from the ILOM including alerts based on pre-set ILOM thresholds.

References: <https://dbatrain.wordpress.com/2011/06/>

[http://docs.oracle.com/cd/E21659\\_01/html/E21660/z40015671004046509.html](http://docs.oracle.com/cd/E21659_01/html/E21660/z40015671004046509.html)

#### NEW QUESTION 30

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 32

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 33

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

#### NEW QUESTION 35

Which two statements are true about the use of DBFS, ACFS, and external file systems and components in an X5 Database Machine environment?

- A. ACFS uses storage in a database on the database machine
- B. DBFS supports file system snapshots.
- C. ACFS supports file system snapshots.
- D. Storage Area Network (SAN) –based file systems may be accessed directly from the storage servers in a database machine.
- E. Oracle Sun ZFS Storage Appliance- based files are not automatically protected by Data Guard.

**Answer:** BC

**Explanation:** DBFS SecureFiles Store Capabilities include support for file system snapshots. Oracle ACFS includes advanced features such as file system snapshot.

References:

<http://www.oracle.com/technetwork/ru/database/dbfs-sf-ooow2009-v2-160969.pdf> <http://www.oracle.com/technetwork/database/database-technologies/cloud-storage/benefits-of-oracle-acfs-2379064.pdf>

#### NEW QUESTION 40

Which three statements are true about Automatic Hard Disk Scrubbing and repair on high- capacity storage servers in an X5 Database Machine?

- A. Disk repair is made immediately from Smart Flash Cache for corruption on disk when the corresponding data is clean in the Smart Flash Cache.
- B. Hard disks are scanned every two weeks by default but only when disks are idle.
- C. Hard disks are scanned every two weeks by default in all situations.
- D. Disk repair is made immediately from Smart Flash Cache for corruption on disk when the corresponding data is dirty in the Smart Flash Cache.
- E. Disk repair may be made from another mirror copy if the corrupt area is part of a Normal or High Redundancy Disk Group.

**Answer:** ACE

**Explanation:** The default schedule of scrubbing is every two weeks.

References:[http://docs.oracle.com/cd/E80920\\_01/DBMMN/maintaining-exadata-storage-servers.htm](http://docs.oracle.com/cd/E80920_01/DBMMN/maintaining-exadata-storage-servers.htm)

#### NEW QUESTION 43

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

#### NEW QUESTION 44

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 49

As part of planning for database consolidation, you grouped your databases into two categories based on different technical and business objectives.

Which three statements are true about possible configurations for your Exadata X6 Database Machine fabric?

- A. The storage grid may be partitioned when deploying a single-rack Database Machine configuration using virtualization.
- B. The database grid may be partitioned when deploying a multirack Database Machine configuration.
- C. A single database cluster benefits from accessing multiple storage grids.
- D. Multiple database clusters may access the same storage grids.
- E. Multiple database clusters never benefit from accessing a single storage grid.

**Answer:** BCD

**NEW QUESTION 53**

Which three storage components are available after nonvirtualized standard deployment of an X5 Database Machine using high-capacity disks?

- A. mirrored system area on hard disk 0, hard disk 1, and hard disk 2
- B. Exadata Smart Flash Cache using all of the flashdisk space
- C. the DBFS\_DG diskgroup with high redundancy
- D. mirrored system area on hard disk 0 and hard disk 1
- E. the RECO\_<DBM\_Name> ASM diskgroup
- F. the DATA\_<DBM\_Name> ASM diskgroup

**Answer:** CEF

**Explanation:** The first two disks of Exadata Storage Server are system disks. Oracle Exadata Storage Server Software system software resides on a portion of each of the system disks. These portions on both system disks are referred to as the system area. The default configuration on Exadata is to have 3 diskgroups - DATA, RECO, and DBFS\_DG.

**NEW QUESTION 54**

You plan to monitor the ASM configuration on an X5 Database Machine as part of your role supporting Exadata-based ASM diskgroups. You want to check for potential space problems that take ASM mirroring requirements into account. Which two values would you monitor from V\$ASM\_DISKGROUP or by using the ASMCMD LSDG command?

- A. cold\_used\_mb
- B. total\_mb
- C. required\_mirror\_free\_mb
- D. free\_mb
- E. usable\_file\_mb

**Answer:** BD

**Explanation:** Determine the Amount of Available Space

To increase the size of the disks in a disk group you must either have unallocated disk space available, or you have to reallocate space currently used by a different disk group.

Example: View the space currently used by the disk groups. SELECT name, total\_mb, free\_mb, total\_mb - free\_mb used\_mb, round(100\*free\_mb/total\_mb,2) pct\_free

FROM v\$asm\_diskgroup ORDER BY 1;

```
SELECT name, total_mb, free_mb, total_mb - free_mb used_mb, round(100*free_mb/total_mb
FROM v$asm_diskgroup
ORDER BY 1;
```

NAME	TOTAL_MB	FREE_MB	USED_MB	PCT_FREE
DATA1	68812800	9985076	58827724	14.51
RECO1	94980480	82594920	12385560	86.96

The example above shows that the DATA1 disk group has only about 15% of free space available while the RECO1 disk group has about 87% free disk space. The PCT\_FREE displayed here is raw free space, not usable free space. Additional space is needed for rebalancing operations.

References:[http://docs.oracle.com/cd/E80920\\_01/SAGUG/exadata-administering-asm.htm#SAGUG20526](http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-administering-asm.htm#SAGUG20526)

**NEW QUESTION 57**

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

**NEW QUESTION 62**

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 64**

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 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

Identify two reasons why the “physical read total bytes” statistic is 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.
- B. The table is an IOT and has an overflow segment, causing “cell multiblock physicalreads” to be requested by the database instance, resulting in additional I/O.
- 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.
- D. The table is an index clustered table, causing “cell single block physical reads” to be requested by the database instance, resulting in additional I/O.
- 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.

**Answer:** BE

**NEW QUESTION 66**

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 68**

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 71

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 saves 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 74

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 75

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 /uo1/stage. You have downloaded the oplan utility to the first database server and run:

```
$ORACLE_HOME/oplan/oplan generateApplySteps /uo1/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 80

Which three are sources for alerts from storage servers in an X5 Database Machine?

- A. Software errors reported by the MS process on the storage servers
- B. Hardware sensor alerts reported by the CELLSRV process on the storage servers
- C. Software errors reported by the RS process on the storage servers
- D. Hardware sensor alerts detected by the ILOM on the storage servers
- E. Metrics that have administrator-defined thresholds on the storage servers

**Answer:** ABE

**Explanation:** AB: 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. In addition, the MS process can generate alerts for important storage cell hardware or software events.

#### NEW QUESTION 81

Examine this list of software components:

1. Oracle Virtual Machine (OVM)
2. Oracle Enterprise Manager Agent (OMA)
3. ASM instance
4. RDBMS instance
5. Automatic Diagnostic Repository Command Interpreter (ADRCI)
6. CELLCLI
7. Cell Server (CELLSRV)
8. diskmon

Identify the location where these software components can run in the standard Exadata Database Machine deployment.

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

**Answer:** B

**Explanation:** Automatic Diagnostic Repository Command Interpreter (ADRCI) can be used on an Exadata storage server.

#### NEW QUESTION 82

Which two options should be used to identify a damaged or failing flash card on an X6 Database Machine high-capacity storage server?

- A. Using the CELLCLI CALIBRATE command on the storage server after logging in as the root user
- B. Using the CELLCLI LIST LUN DETAIL command as the celladmin user
- C. Using the CELLCLI CALIBRATE command on the storage server after logging in as the celladmin
- D. Hardware monitoring using the storage server ILOM

**Answer:** AD

**Explanation:** The CALIBRATE command runs raw performance tests on cell disks, enabling you to verify the disk performance before the cell is put online. You must be logged on to the cell as the root user to run CALIBRATE.

Oracle ILOM automatically detects system hardware faults and environmental conditions

on the server. If a problem occurs on the server, Oracle ILOM will automatically do the following:

Illuminate the Service Required status indicator (LED) on the server front and back panels. Identify the faulted component in the Open Problems table.

Record system information about the faulted component or condition in the event log.

References:[http://docs.oracle.com/cd/E80920\\_01/SAGUG/exadata-storage-server-cellcli.htm#SAGUG20617](http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-storage-server-cellcli.htm#SAGUG20617)

#### NEW QUESTION 84

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

**NEW QUESTION 87**

You installed ASR Manager on a stand-alone server and configured Auto Service Request (ASR) for your X5 Database Machine and its assets. Which three statements are true about this configuration?

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

**Answer:** BCF

**Explanation:** B: Oracle ASR Manager only processes SNMP traps that are sent from IP addresses that Oracle ASR Manager recognizes.

Example of Exadata Storage Server SNMP Trap

This example shows the SNMP trap for an Exadata Storage Server disk failure. The corresponding hardware alert code has been highlighted.

2011-09-07 10:59:54 server1.example.com [UDP: [192.85.884.156]:61945]:

RFC1213-MIB::sysUpTime.0 = Timeticks: (52455631) 6 days, 1:42:36.31 SNMPv2-SMI::snmpModules.1.1.4.1.0 = OID: SUN-HW-TRAP-MIB::sunHwTrapHardDriveFault

SUN-HW-TRAP-MIB::sunHwTrapSystemIdentifier = STRING: Sun Oracle Database Machine

Etc.

C (not D): The ASR Manager uses the SNMP GET protocol to query ASR assets for additional fault information.

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

References: [http://docs.oracle.com/cd/E80920\\_01/ASXQI/toc.htm](http://docs.oracle.com/cd/E80920_01/ASXQI/toc.htm)

[https://docs.oracle.com/cd/E37710\\_01/install.41/e18475/ch5\\_troubleshooting.htm#ASRUD\\_331](https://docs.oracle.com/cd/E37710_01/install.41/e18475/ch5_troubleshooting.htm#ASRUD_331)

**NEW QUESTION 91**

You are in the process of upgrading your X5 Database Machine half rack to a full rack. The new storage servers are called DM01CEL08 through DM01CEL14. After creating 96 new griddisks, you issued this SQL statement:

```
SQL> ALTER DISKGROUP DATA ADD DISK
2> 'O/*/DATA*DM01CEL08*'
3> 'O/*/DATA*DM01CEL09*'
4> 'O/*/DATA*DM01CEL10*'
5> 'O/*/DATA*DM01CEL11*'
6> 'O/*/DATA*DM01CEL12*'
7> 'O/*/DATA*DM01CEL13*'
8> 'O/*/DATA*DM01CEL14*'
9> REBALANCE POWER 512;
```

How many failgroups if any, will be added to the DATA diskgroup by executing this SQL statement?

- A. 0, because the new griddisks will be added to the existing failgroups
- B. 12, consisting of seven griddisks each
- C. 96, consisting of seven griddisk each
- D. 7, consisting of 12 griddisks each
- E. 1, consisting of all 96 griddisks

**Answer:** D

**Explanation:** The number of failure groups equals the number of Exadata Cells. Each failure group is composed of a subset of grid disks in the Oracle ASM disk group that belong to a single storage cell.

References: [http://docs.oracle.com/cd/E80920\\_01/SAGUG/exadata-administering-asm.htm](http://docs.oracle.com/cd/E80920_01/SAGUG/exadata-administering-asm.htm)

**NEW QUESTION 95**

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

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

2nd - Questions and Answers in PDF Format

### **1z0-070 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

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The 1z0-070 Practice Test Here](#)**