



1Z0-070^{Q&As}

Oracle Exadata X5 Administration

Pass Oracle 1Z0-070 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass4itsure.com/1z0-070.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers





QUESTION 1

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

Correct 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
RECOC1	94980480	82594920	12385560	86.96

The example above shows that the DATA1 disk group has only about 15% of free space available while the RECOC1 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-administeringasm.htm#SAGUG20526



QUESTION 2

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

Correct 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-9FF8B5B0-3481-4B73-89D3108CBD7EB989.htm#ELMOG76346 http://docs.oracle.com/cd/E80920_01/DBMIN/configuring-exadata.htm#DBMIN21274

QUESTION 3

Which three statements are true about the EXADCLI utility?

- A. It may be run interactively.



- B. It can be used to execute EXACLI scripts on multiple storage servers in parallel.
- C. It can be used to execute O/S commands on multiple storage servers in parallel.
- D. It can be used to execute EXACLI commands on multiple storage servers in parallel.
- E. It uses the same security mechanism as the EXACLI command.
- F. It may be used to execute DBMCLI commands on multiple database nodes in parallel.

Correct Answer: DEF

Explanation:

The exadcli utility runs commands on multiple remote nodes in parallel threads.

You can issue an ExaCLI command to be run on multiple remote nodes. Remote nodes are referenced by their host name or IP address. Unlike dcli, exadcli can only execute ExaCLI commands. Other commands, for example, shell commands, cannot be executed using exadcli.

Incorrect Answers:

A: Note that exadcli runs ExaCLI in a “no-prompt” mode. This means that if user interaction is needed (for example, if you need to enter a password or if you need to verify that the certificate from a remote node is valid), then exadcli will exit with an error.

References:

http://docs.oracle.com/cd/E80920_01/DBMMN/exadcli.htm#DBMMN-GUID-4AE469A6-F291-4737-B975-F1B4B91D0BA0
https://docs.oracle.com/cd/E62172_01/html/E63692/z400007d1478481.html

http://docs.oracle.com/cd/E80920_01/DBMMN/exadcli.htm#DBMMN-GUID-4AE469A6-F291-4737-B975-F1B4B91D0BA0
https://docs.oracle.com/cd/E62172_01/html/E63692/z400007d1478481.html

QUESTION 4

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.

Correct Answer: D

**QUESTION 5**

Which four statements are true concerning the configuration or use of Enterprise Manager Cloud Control to monitor and manage Exadata Database Machine components?

- A. Computer nodes forward their SNMP traps to the Management Server process on the same computer node.
- B. Storage nodes forward their SNMP traps to the Management Server process on the same storage node.
- C. Cloud control can monitor and manage a Database Machine compute node expansion rack.
- D. Oracle management agents must only be installed on one storage node in a Database Machine to enable monitoring and management of all storage nodes.
- E. Cloud Control cannot monitor partitioned Exadata Database Machine systems.
- F. Cloud Control can monitor and manage a Database Machine storage Expansion rack.
- G. Oracle management agents must only be installed on one compute node in a Database Machine to enable monitoring and management of all compute nodes and storage nodes.

Correct Answer: BCFG

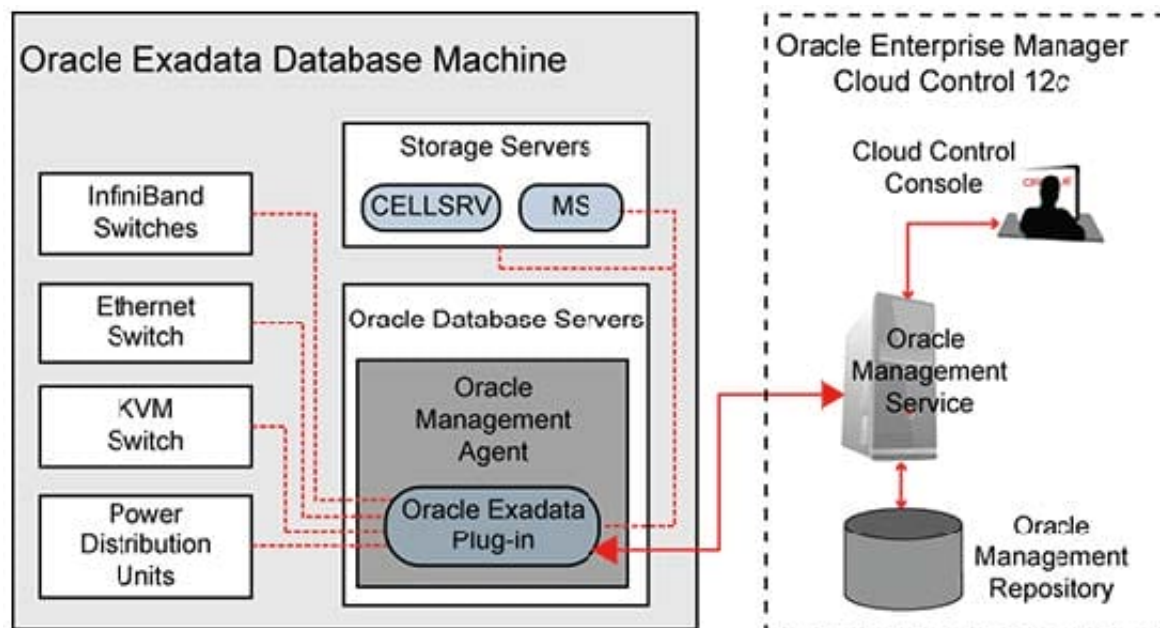
Explanation:

B (not A): Oracle Exadata components—including database and storage servers, switches, and power distribution units (PDUs)—use SNMP to raise alerts and report monitoring information.

Oracle's CellCLI interface is a Java-based framework delivered by the storage cell's management server process (MS) that provides administrative capabilities to your storage server entities.

C: Cloud control can monitor the Exadata Database Machine.

Figure: Oracle Enterprise Manager Cloud Control 12c monitoring architecture. Note the Management Server process (MS) on the Storage servers.



G (not D): Before using Oracle Enterprise Manager Cloud Control 12c with Oracle Exadata, an Oracle Management Agent and Oracle Exadata plug-in must be installed on every Oracle Exadata database server (compute node). This agent monitors software targets, such as the database instances and Oracle Clusterware resources, on the database servers. The plug-in enables monitoring of other hardware components in Oracle Exadata, including the storage servers, switches, and power distribution units.

References: <http://www.oracle.com/technetwork/articles/servers-storage-admin/monitor-exadata-em122291964.html>

QUESTION 6

You plan to migrate a very large database supporting a DSS workload to your new X5 Database Machine.

It will be the only database on this full rack.

Which three statements are true about Database Machine features that improve performance for the DSS workload?

- A. Smart Storage operations can improve the performance of joins.
- B. Smart Storage operations can improve the performance of scans.
- C. Hybrid Columnar Decompression overheads can be offloaded from the database servers for index full scans.
- D. Full table scan operations can improve due to the default Smart Flash Cache implementation.
- E. Hybrid Columnar Compression can reduce the amount of physical I/O required to scan large tables.

Correct Answer: BCD

Reference: <http://www.informit.com/articles/article.aspx?p=2418151&seqNum=3>

Reference: <http://www.informit.com/articles/article.aspx?p=2418151andseqNum=3>

**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 filesystembased repository.
- E. Cell alerts are written to disk every hour by default.

Correct 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>

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

QUESTION 8

Identify three valid configuration options that can be implemented during the initial configuration process of a new X5 Database Machine by using the Exadata Deployment Assistant.

- A. Virtualized Compute Node configuration
- B. OS users that own the Grid Infrastructure installation
- C. Cluster configuration for up to 32 clusters
- D. The size for the system area on the first two disks in each cell
- E. Virtualized Storage node configuration



Correct Answer: ACD

Explanation:

Oracle Exadata Deployment Assistant has two main phases

1.
the customer provides naming, DNS, NTP, and networking details. This information is collected using the Oracle Exadata Deployment Assistant user interface. A configuration file is generated

2.
the configuration file is pushed to the database node, and the validation and installation is

Oracle Exadata Deployment Assistant can configure Oracle Exadata Database Machine, Recovery Appliance, or Oracle SuperCluster.

References:

https://docs.oracle.com/cd/E55822_01/AMOGD/configurator.htm#AMOGD22056

<https://www.toadworld.com/platforms/oracle/w/wiki/11469.oracle-exadata-deployment-assistance-oeda>

QUESTION 9

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

Correct 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.

QUESTION 10

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.

Correct 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-GUID42945059-13FD-4F6A-B7FA-A1201D16238F http://docs.oracle.com/cd/E80920_01/DBMSO/exadata-whatsnew.htm#DBMSO22120

QUESTION 11

Which two network switch failure scenarios on a standalone Exadata Database Machine X5-2 Full Rack will affect database availability?

- A. failure of the Ethernet switch only
- B. failure of both the InfiniBand leaf switches
- C. failure of a single InfiniBand leaf switch and the Ethernet switch
- D. failure of a single InfiniBand leaf switch
- E. failure of both the InfiniBand leaf switches and the Ethernet switch

Correct Answer: BE

Explanation:

Ethernet switch for administrative connectivity to servers in the Database Machine

Connect any combination of up to 18 Exadata Database Machine racks or Exadata Storage Expansion

Racks via the InfiniBand fabric. Larger Configurations can be built with external InfiniBand switches.

Connected racks can be any combination of v2, X2, X3 or X4 generation hardware.



Reference: <https://www.oracle.com/technetwork/database/exadata/exadata-storage-expansion-x5-2-ds-2406252.pdf>

Reference: <https://www.oracle.com/technetwork/database/exadata/exadata-storage-expansion-x5-2-ds-2406252.pdf>

QUESTION 12

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.

Correct 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

https://docs.oracle.com/cd/E37710_01/install.41/e18475/ch5_troubleshooting.htm#ASRUD331

QUESTION 13

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 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

Correct Answer: D

Explanation: CT_IO_RQ_SM The cumulative number of small I/O requests issued by the category for hard disks. A



large value indicates a heavy I/O workload from this category. CT_IO_RQ_LG The cumulative number of large I/O requests issued by the category for hard disks. A large value indicates a heavy I/O workload from this category. CT_IO_WT_SM_RQ The average IORM wait time per request for small I/O requests issued to hard disks by an IORM category. CT_IO_WT_LG_RQ The average IORM wait time per request for large I/O requests issued to hard disks by an IORM category.

Incorrect Answers:

A: CT_IO_RQ_SM_SEC

This metric is derived from CT_IO_RQ_SM. It specifies the rate of small I/O requests issued by the category for hard disks. Its units are number of I/O requests per second. A large value indicates a heavy I/O workload from this category in the past minute.

B: CT_IO_UTIL_SM

The percentage of disk resources utilized by small requests from this category.

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

QUESTION 14

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: 11.2.2.3.1.110429.1Active image activated: 2011-05-09 16:26:36 -0400Active image status:
successActive image partition on device: /dev/md6Active software partition on device: /dev/md8
```

```
In partition rollback: Impossible
```

```
Cell boot usb partition: /dev/sdm1Cell boot usb version: 11.2.2.3.1.110429.1
```

```
Inactive image version: 11.2.1.2.6Inactive image activated: 2010-10-04 23:59:16 -0400Inactive image status:
successInactive system partition on device: /dev/md5Inactive 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 physicaldisk 1
- B. /dev/md6 and /dev/md8 are located only on physicaldisk 1
- 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

Correct Answer: BC

**QUESTION 15**

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.

Correct 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

[1Z0-070 PDF Dumps](#)

[1Z0-070 Practice Test](#)

[1Z0-070 Braindumps](#)