



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.passapply.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 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.

Correct 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-exadatapart-3-migrating-databases-to-exadata-database-machine-best-practices>

QUESTION 2

Which three statements are true about Recovery Manager (RMAN) daily differential incremental backup strategies on an X5 Database Machine for a database having 25% or more of its blocks modified each day and which has an 8 k block size?

- A. Fast incremental backups when 50% or more of the blocks have changed since the last backup, will run as slowly as normal incremental backup.
- B. Enabling Block Change Tracking (BCT) on the database can result in reduced consumption of storage network bandwidth.
- C. Enabling Block Change Tracking (BCT) on the database can result in a reduction of physical I/O on the cells during incremental backups.
- D. For level-1 backups, Block Change Tracking (BCT) is most beneficial when more than 25 percent of the blocks have changed since the last backup.
- E. For level-0 backups, Block Change Tracking (BCT) is most beneficial when more than 25 percent of the blocks have changed since the last backup.



F. cellsrv returns only blocks that have changed since the last backup.

Correct Answer: ACF

Explanation:

A: Fast Incremental backups is possible with Block change tracking, which is initially introduced from version 10.2 onwards, by this tool it's very useful to reduce the RMAN incremental backup duration. If the changes are something around 20% then in this situation BCT helps a lot.

C: Exadata Storage Server offload capability combined with RMAN block change tracking will efficiently perform large I/Os at the storage-tier level, returning only individual changed blocks for incremental backups and increasing the backup performance of the system.

Note: Level 1 backup: A level 1 backup includes only those blocks that have been changed since the "parent" backup was taken. Remember a parent backup may be either a level 0 or a level 1 backup. Block change tracking allows indeed the highest benefit for databases where the changes are not so high,

Level 0 backup: A level 0 incremental backup is physically identical to a full backup and it includes every data block in the file except empty blocks. The only difference is that the level 0 backup is recorded as an incremental backup in the RMAN repository, so it can be used as the parent for a level 1 backup.

References: http://www.dba-oracle.com/t_rman_backup_types.htm
<http://www.oracle.com/technetwork/database/availability/maa-tech-wp-sundbm-backup-11202-183503.pdf>
<https://www.toadworld.com/platforms/oracle/w/wiki/11124.fast-incremental-backups-active-data-guard>

QUESTION 3

You plan to consolidate multiple mission-critical databases onto a single rack Exadata X6 database Machine.

You do not plan to use virtualization on the database servers.

You have identified two database categories, based on differing availability and maintenance objectives.

Which two choices will be the best solution?

- A. Create a single database cluster that accesses a single storage grid.
- B. Create a single database cluster that accesses multiple storage grids.
- C. Partition the Exadata storage servers into two separate storage grids.
- D. Create a single storage grid and share space on the cell disks using Exadata storage realms.
- E. Partition the database servers into two separate clusters.

Correct Answer: BE

QUESTION 4

You are planning the monitoring configuration for your X5 Database Machine.



Which two components are monitored directly through the use of Exadata–specific Enterprise Manager Plug-Ins?

- A. the database server Clusterware
- B. the storage server ILOM
- C. ASM instances
- D. the database server O/S
- E. the Infiniband switches

Correct Answer: BE

Explanation:

The Oracle Enterprise Manager Grid Control Exadata Monitoring plug-in bundle allows you to monitor the following key components of Exadata machine:

1.

Oracle ILOM: The plug-in monitors the Oracle ILOM card in a database server for hardware events and records sensor data.

2.

InfiniBand Switch: The plug-in enables Enterprise Manager Grid Control to monitor the Oracle DataCenter36 Infiniband Switch.

3.

Avocent MergePoint Unity Switch: The Plug-in enables Enterprise Manager Grid Control to monitor KVM (keyboard, video or visual display unit, mouse) targets. The plug-in provides status of the KVM and the event occurrences like Factory Defaults Set, Fan Failure, Aggregated TargetDevice Status, Power Supply Failure, Power Supply Restored, Reboot Started, Temperature Out of Range on the KVM target.

4.

Cisco Switch: The plug-in enables Enterprise Manager Grid Control to monitor Cisco Switch targets. This plug-in will monitor the Cisco switch's CPU, memory, temperature, network interfaces, system information, fan, and power supply metrics.

5.

Power Distribution Unit: The plug-in will be used to monitor PDU's actual current value of Phase1, Phase2, and Phase3

References: <http://www.oracle.com/technetwork/oem/grid-control/exadata-plug-in-bundle-188771.html>

QUESTION 5

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.

Correct Answer: BCE

QUESTION 6

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

- 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 clusters

Correct 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 cell disk. In a standard Exadata deployment we

create griddisks on hard disk based cell disks 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:



1.

Add the storage cell to the Exadata Cell realm

2.

Configure security on the Oracle Exadata Storage Server grid disks

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

QUESTION 7

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

Correct Answer: ABD

QUESTION 8

Identify three facilities that can use X5 Extreme Flash Exadata Storage Server flashdisk space.

- A. Cell System Area
- B. Exadata Smart Flash Cache
- C. Database Smart Flash Cache
- D. cell boot flash image
- E. Exadata Smart Flash Log

Correct Answer: BDE

QUESTION 9



In which three situations will a Smart Scan occur on a table for which Smart Scan is possible, assuming that nothing in the environment prevents a Smart Scan?

- A. The ASM diskgroup containing the table's tablespace has a 4 MByte AU size.
- B. Executing queries in parallel.
- C. setting `_serial_direct_read= true` in the session issuing the SQL statements and executing a serial query.
- D. Having direct path reads used at execution time.
- E. An update to the table is made identifying the row to be modified by primary key.
- F. An index range scan occurs for the table.

Correct Answer: BCD

Explanation: Whether you get a direct path read and smart scan, depends on the current buffer cache size, how big segment you're about to scan and how much of that segment is actually cached at the moment. This dynamic decision unfortunately can cause unexpected surprises and variance in your report/batch job runtimes.

To work around these problems and force a direct path read/smart scan, you can either:

1.

(B) Run your query in parallel as parallel full segment scans will use direct path reads, unless your `parallel_degree_policy = AUTO`, then you may still get buffered reads thanks to the dynamic in-memory parallel execution decision of Oracle 11.2

2.

(C) Run your query in serial, but force the serial direct path reads by setting `_serial_direct_read = TRUE` (or `ALWAYS` in 11.2.0.2+)

D: One of the most common Exadata performance problems is that the direct path reads (and thus also Smart Scans) don't sometimes kick in when running full scans in serial sessions.

Note: Smart Scan is a subset of Cell Offload Processing and is perhaps the feature that is most highly associated with Exadata. Smart Scan allows Oracle databases to push query information down to the storage level, specifically:

1.

Filtration (predicate information)

2.

Column Projection

3.

Join Filtering

Incorrect Answers:

F: Objects such as clusters and index organized tables cannot be Smart Scanned.

References: http://blog.tanelpoder.com/2013/05/29/forcing-smart-scans-on-exadata-is-_serial_direct_readparameter-



safe-to-use-in-production/

QUESTION 10

Which two statements are true about the EXACLI utility?

- A. It can execute commands on multiple cells in parallel.
- B. It supports all CELLCLI commands.
- C. It requires the same privileges as the celladmin linux user to create CELLDISKS and GRIDDISKS.
- D. It can be executed using the EXADCLI utility.
- E. It can be executed directly on cells.
- F. It requires access through a cell user that has role-based privileges granted.

Correct Answer: AD

Explanation:

The exadcli utility facilitates centralized management across an Oracle Exadata system by automating the execution of ExaCLI commands.

Exadcli executes the commands on a set of remote nodes and returns the output to the centralized management location where the exadcli utility was run.

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

Reference: <https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmmn/exadcli.html#GUID-1C738F05-2A69-4B75-BB1E-B578C9081487>

Reference: <https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmmn/exadcli.html#GUID-1C738F05-2A69-4B75-BB1E-B578C9081487>

QUESTION 11

Which two statements are true about Exadata storage server and database server alerts on an X5 Database Machine?

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



Correct Answer: BC

Reference:

<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/sagug/exadata-storage-server-software-introduction.html#GUID-3E48425A-AB8A-4E62-80C4-BACA65A1F8D3>

https://docs.oracle.com/cd/E91266_01/EMXIG/GUID-FB58204F-2D97-41BC-9AA7-10BFF920B5B4.htm#EMXIG304

<https://docs.oracle.com/en/engineered-systems/exadata-database-machine/sagug/exadata-storage-server-software-introduction.html#GUID-3E48425A-AB8A-4E62-80C4-BACA65A1F8D3>

https://docs.oracle.com/cd/E91266_01/EMXIG/GUID-FB58204F-2D97-41BC-9AA7-10BFF920B5B4.htm#EMXIG304

QUESTION 12

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



QUESTION 13

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 14

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_ ASM diskgroup



F. the DATA_ ASM diskgroup

Correct 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.

QUESTION 15

Which three are true concerning Smart Scans?

- A. The decision to perform a Smart Scan is made by the Optimizer.
- B. Smart Scans are only possible during full table scans or fast full index scans.
- C. Smart Scans are possible on Index Organized Tables (IOTs).
- D. Smart Scans always outperform scans that are not offloaded.
- E. Smart Scans are possible on heap organized tables.
- F. Smart Scans are possible on uncompressed B*Tree indexes.

Correct Answer: ABF

Explanation:

B: In order for queries to take advantage of Exadata's Offloading capabilities, the optimizer must decide to execute a statement with a Full Table Scan or a Fast Full Index Scan.

F: Smart Scans can be used for full scanning through B*Tree index segments.

References: Expert Oracle Exadata - K. Osborne, et al., (Apress, 2011) BBS, page 346

[Latest 1Z0-070 Dumps](#)

[1Z0-070 VCE Dumps](#)

[1Z0-070 Practice Test](#)