

Oracle 1Z0-070 Exam

Volume: 90 Questions

Question: 1

You plan to migrate a database supporting an OLTP workload to your new X5 Database Machine.

The current database instance supports a large number of short duration sessions and a very high volume of short transactions.

Which three X5 Database Machine features can improve performance for this type of workload?

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

Answer: A,C,D

Question: 2

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

Oracle 1Z0-070 Exam

- 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

Question: 3

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

Which two components need to be monitored?

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

Answer: C,D

Question: 4

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

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

Oracle 1Z0-070 Exam

the same cluster must run the same version of Oracle Virtual Machine (OVM).

Answer: D

Question: 5

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: B,C,F

Question: 6

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.

Oracle 1Z0-070 Exam

F. Smart Scans are possible on uncompressed B*Tree indexes.

Answer: A,B,F

Question: 7

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: A,E

Question: 8

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

A. Backups are fast because the Infiniband network is faster than fiber channel storage, so RMAN read I/O operations are faster than in traditional storage networks.

B. Backups are fast, because physical I/Os operations are reduced by the use of the Storage Index by cellsvr.

C. Backups are fast, because data blocks are read from the Database Flash Cache.

D. Backups are fast, because all data blocks are read from the Exadata Smart Flash Cache.

E. For incremental backups using Block Change Tracking, cellsvr filters blocks and returns only those that have changed since the last backup.

Answer: D,E

Oracle 1Z0-070 Exam

Question: 9

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

Question: 10

Which three are recommended options on an X5 Exadata Database machine, for staging files that will be loaded as external tables?

- A. on an ACFS file system stored in a staging database on the Database Machine
- B. a ZFS over NFS file system mounted on a database server where the external table will be accessed
- C. an Exadata-based ACFS file system on the Database Machine
- D. a DBFS file system stored in a staging database on the Database Machine
- E. on local storage on one or more storage servers that are accessible to the database server where the load will be performed

Oracle 1Z0-070 Exam

Answer: A,B,E

Question: 11

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: A,D

Question: 12

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

```
alter iormplan objective = off
```

```
alter iormplan active
```

There are no database or category plans defined.

You are suffering I/O performance problems at certain times, which vary by day and week. DSS and batch workloads perform well, but your OLTP workloads suffer poor response times when running at the same time as these other workloads.

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

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

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

Answer: A

Oracle 1Z0-070 Exam

Question: 13

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

```
SQL> SELECT s.name, m.value/1024/1024 MB FROM V$SYSSTAT s, V$MYSTAT m
2 WHERE s.statistic# = m.statistic# AND
3 (s.name LIKE 'physical % total bytes' OR s.name LIKE 'cell phys%'
4 OR s.name LIKE 'CELL 10%');
```

NAME	MB
physical read total bytes	19047.2266
physical write total bytes	0
cell physical IO interconnect bytes	4808.85828
cell physical IO bytes pushed back due to excessive CPU on cell	0
cell physical IO bytes saved during optimized file creation	0
cell physical IO bytes saved during optimized RMAN file restore	0
cell physical IO bytes eligible for predicate offload	18005.6953
cell physical IO bytes saved by storage index	0
cell physical IO interconnect bytes returned by smart scan	3767.32703
cell IO uncompressed bytes	18005.6953

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

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

Answer: C,D

Question: 14

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.

Oracle 1Z0-070 Exam

- 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: B,C,D

Question: 15

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 on the first database server and is located at /u01/stage.

You have downloaded the oplan utility to the first database server and run:

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

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

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

Answer: A,C

Oracle 1Z0-070 Exam

Question: 16

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: A,B,E

Question: 17

Which three are true about Exadata Smart Flash Log?

- A. Databases on the Database Machine use Exadata Smart Flash Log by default.
- B. I/O Resource Manager database plans can be used to enable or disable Exadata Smart Flash Log for individual databases.
- C. LGWR will not wait for writes to Exadata Smart Flash Log if the write to a disk-based logfile completes first.
- D. I/O Resource Manager category plans can be used to enable or disable Exadata Smart Flash Log for different I/O categories.
- E. The use of Exadata Smart Flash Logs is mandatory for support of production databases.

Answer: A,B,C

Question: 18

You plan to migrate an Oracle database supporting an online transaction processing (OLTP) workload to an X5 Exadata Database Machine.

The database machine is running database version 12.1.0.2.

The source database has these attributes:

Database version: 12.1.0

Byte order: Big Endian

Database size: 24 TB

Oracle 1Z0-070 Exam

You have decided to perform physical migration using the Transportable Tablespace method. Examine this list of possible steps to accomplish this task.

1. Create a new 12.1.0.2 Exadata –based database.
2. Export the source system metadata using transportable Tablespace (TTS).
3. Transfer the files to the Database Machine and use the RMAN CONVERT command to change them to the little endian format.
4. Import the metadata into the target database.
5. Use the RMAN CONVERT to change the database files to the little endian format on the source system.

Identify the required steps in the correct order.

- A. 1, 2, 4, 5
- B. 1, 2, 5, 3, 4
- C. 1, 2, 3, 4, 5
- D. 1, 2, 3, 4
- E. 1, 2, 5, 4, 3
- F. 1, 2, 5, 4

Answer: D

Question: 19

An Exadata storage server physical disk on an X5 high-capacity full rack entered the predictive failure state.

Identify the two steps that you must perform to replace this failed physical disk.

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