

PDFVCE



Choose the version that fits your needs	PDF Version	Desktop Test Engine	Online Test Engine
Latest and Up-to-Date exam dumps with real exam questions answers.	✓	✓	✓
Get 12-Months free updates without any extra charges.	✓	✓	✓
Experience same exam environment before appearing in the certification exam.	✗	✓	✓
100% exam passing guarantee in the first attempt.	✓	✓	✓
20% discount on more than one license and 30% discount on 5+ license purchases.	✗	✓	✓
100% secure purchase on SSL.	✓	✓	✓
Completely private purchase without sharing your personal info with anyone.	✓	✓	✓

<http://www.pdfvce.com>

Highly Efficiently Exam Tool and Effective Exam Practice Materials

Exam : **1z0-070**

Title : Oracle Exadata X5
Administration

Vendor : Oracle

Version : DEMO

NO.1 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.** A single database cluster benefits from accessing multiple storage grids.
- C.** The database grid may be partitioned when deploying a multirack Database Machine configuration.
- D.** Multiple database clusters may access the same storage grid.
- E.** Multiple database clusters never benefit from accessing a single storage grid.

Answer: B,C,D

NO.2 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.** 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.
- B.** 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.
- C.** Create a database plan using EXADCLI calling EXACLI to give OLTP I/O the highest priority from all six databases.
- D.** 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.

Answer: B

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

NO.3 Which two statements are true regarding the use of Auto Service Request (ASR) with an X6 Database Machine?

- A.** The database server MS process must have SNMP traps configured to use the management network for notifications to ASR Manager.

B. The database server ILOMs must use SMTP over the management network for notifications to ASR Manager.

C. The database server ILOMs must have SNMP traps configured to use the management network for notifications to ASR Manager.

D. The storage server ILOMs must have SNMP traps configured to use the management network for notifications to ASR Manager.

Answer: C,D

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

NO.4 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 flash cache compression technique

B. Superior compression capability designed specifically for archival data

C. Superior compression capability designed specifically for OLTP workloads

D. Cell offload processing for indexed-organized table access

E. Superior compression capability designed specifically for data warehouse tables

F. Columnar storage capability for data in row major data blocks that is held in flash cache

Answer: B,C,E

Explanation:

There are 6 different kinds of Table Data Compression methods:

NO.5 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 range scans on the corresponding table.
- C. if Smart only occur instead of index skip scans on the corresponding table.
- D. if Smart Scans performs equally well to any type of index scan on the corresponding table.

Answer: A,B

NO.6 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. 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.
- B. cellsrv returns only blocks that have changed since the last backup.
- 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. Fast incremental backups when 50% or more of the blocks have changed since the last backup, will run as slowly as normal incremental backup.
- F. Enabling Block Change Tracking (BCT) on the database can result in reduced consumption of storage network bandwidth.

Answer: B,C,E

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>

NO.7 Which three statements are true about bulk data loading capabilities in an X5 Database Machine environment?

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

Answer: A,C,D

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

NO.8 Which two completely prevent a Smart Scan from occurring?

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

Answer: C,E

Explanation:

B: Migrated Rows is a special case of chained rows. Migrated rows still affect performance, as they do in conventional storage situations, but with the additional overhead of reducing the beneficial effects of "Smart Scan" in addition to increasing the number of I/Os E: Smart Scans - broadly speaking and ignoring edge cases - can only transport a maximum of 254 columns from a single (non-HCC) segment. Requesting more columns will simply disable Smart Scans for that segment.

An interesting limitation to Exadata Smart Scans - if more than 254 columns from a table (not HCC compressed, more on that in moment) need to be projected, Smart Scans for that particular segment will be disabled and Exadata will fall back to conventional I/O. This means that the number of columns in the projection clause can make a significant difference to performance, since only Smart Scans allow taking advantage of offloading and particularly avoiding I/O via Storage Indexes.

NO.9 Which two statements are true about data paths used during RMAN backup and restore

operations when media management servers use InfiniBand to connect to an X5 Database Machine?

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

Answer: A,E

Explanation:

The fastest database backup is achieved via InfiniBand.

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

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

References:

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

NO.10 You issued these commands to all Exadata Storage Servers in an X6 Exadata Database Machine using dcli:

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

There are no database or category plans defined.

You are encountering disk I/O performance problems at certain times, which vary by day and week. DSS and Batch workloads perform well some of the time.

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

You wish to have disk I/O managed so that performance will be optimized for all workloads.

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

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

Answer: A

Explanation:

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

References:

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

