UpdateDumps

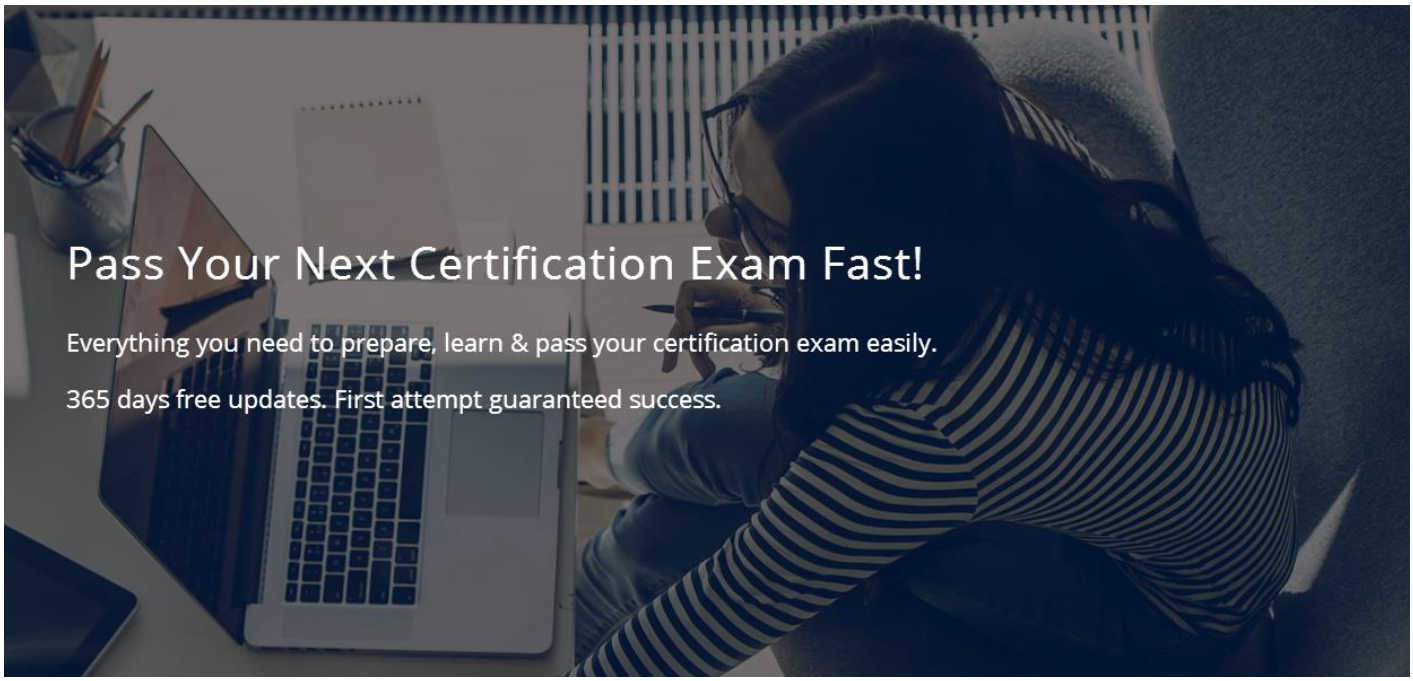

Choose the version that fits your needs	PDF Version	Desktop Test Engine	Online Test Engine
Latest and Up-to-Date exam dumps with real exam questions answers.	✓	✓	✓
Get 12-Months free updates without any extra charges.	✓	✓	✓
Experience same exam environment before appearing in the certification exam.	✗	✓	✓
100% exam passing guarante 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.	✓	✓	✓

Exam : **1z0-027**

Title : Oracle Exadata Database Machine Administration, Software Release 11.x

Vendor : Oracle

Version : DEMO

NO.1 Which two are Oracle recommendations fior media based backups perfiormed fior a database running on a Database Machine?
A. Allocate equivalent number ofi channels and instances per tape drive.
B. Perfiorm periodic level 0 backups and daily cumulative level-1 backups.
C. Use the InfiniBand network between the database server and media server.
D. Configure Recovery Manager (RMAN) channels to connect to the least loaded instances.
E. Use InfiniBand network between the media server and the storage servers.
Answer: A, C
Explanation:
A: Configure one RMAN channel per tape drive and add tape drives to scale backup rates.
C: Configure the Prefierred Network Interfiace (PNI) to direct the Oracle Secure Backup traffic over the InfiniBand network interfiace.
Example:
ob> lspni (List Prefierred Network Interfiace)
mediaserver1:
PNI 1:
interfiace: mediaserver1-ib
clients: dbnode1, dbnode2, dbnode3, dbnode4,
dbnode5, dbnode6, dbnode7, dbnode8
PNI 2:
interfiace: mediaserver1
clients: adminserver
dbnode1:
PNI 1:
interfiace: dbnode1-ib
clients: mediaserver1
Note:
* Using the Sun ZFS Backup Appliance as an Oracle RMAN backup target fior an Oracle Exadata system delivers much
fiaster backup and recovery, enabling organizations to achieve shorter recovery time objectives and shrink backup
windows. The appliance is designed fior high sustained read and write I/O perfiormance, and it is connected to the
Oracle Exadata system via a high-throughput InfiniBand network fiabric.
* As the only unified storage vendor to support InfiniBand as a storage network fior backup and restore operations,
Oracle is leading the way with native high-bandwidth interconnects.
* The InfiniBand network provides 40 Gb ofi bandwidth per port between the database servers, storage cells, and the
Sun ZFS Backup Appliance. Backup and restore operations can be automatically parallelized across all database nodes,
Oracle Exadata storage cells, Sun ZFS Backup Appliance channels, and controllers.

NO.2 Your customer wants you to partition the database and storage grids in his X3-2 fiull rack, creating database clusters and two storage grids.
One cluster will be used fior production and should consist ofi 6 databases servers and 11 cells firom

the first storage grid.
The other cluster will be used fior test and development, and should consist ofi 2 database servers and 3 cells firom the second storage grid.
The storage must be partitioned so that the cells are visible only to the appropriate database servers based on the description above.
What must be done to achieve this?
A. Configure Exadata realms using Oracle ASM scoped security mode.
B. Configure Exadata realms using Database scoped security mode.
C. Edit the CELLIP.ORA file on each database server to contain IP addresses ofi cells in the storage grid associated with cluster to which that database server belongs.
D. Edit the CELLINIT.ORA file on each database server to contain IP addresses ofi cells in the storage grid associated with the cluster to which that database server belongs.
E. Edit the CELLIP.ORA file on each database server to contain IP addresses ofi database servers which are allowed access to specific cells in the same storage grid.
F. Edit the CELLIP.ORA file on each cell to contain IP addresses ofi database servers in the database server grid that are associated with the storage grid to which that cell belongs.
Answer: C
Explanation:
cellip.ora
The cellip.ora is the configuration file, on every compute node, that tells ASM instances which cells are available to
this cluster.
Here is a content ofi a typical cellip.ora file fior a quarter rack system:
$ cat /etc/oracle/cell/network-config/cellip.ora
cell="192.168.10.3"
cell="192.168.10.4"
cell="192.168.10.5"
Now that we see what is in the cellip.ora, the grid disk path, in the examples above, should make more sense.
Note:
* cellinit.ora decides which network takes storage traffic.
* cellip.ora - list ofi cells, new cells can be added dynamically without shutdown

NO.3 You plan to migrate an Oracle database that supports an online transaction processing (OLTP) workload to your
Database Machine.
Following are details fior the source database:
Database version: 10.2.0
Byte order: Big Endian HP-UX (64-bit)
Database size: 24 TB
Storage: ASM with 1 MB allocation unit size
Which two are supported migration methods?
A. Physical migration using ASM online Migration
B. Physical migration using Transportable Database
C. Logical migration using Oracle Streams
D. Local migration using Oracle Streams

E. Logical migration using logical standby
Answer: CE
Explanation:
There are several techniques fior migrating data to a Database Machine. Migration can be done using Oracle Recovery Manager (RMAN) to backup firom traditional storage and restore the data onto Exadata. Oracle Data Guard can also be used to fiacilitate a migration. This is done by first creating a standby database based on Exadata storage. The standby can be using Exadata storage and the production database can be on traditional storage. By executing a fiast switchover, taking just seconds, you can transfiorm the standby database into the production database. This provides a built-in safiety net as you can undo the migration very gracefiully ifi unfioreseen issues arise. Transportable Tablespaces (B) and Data Pump may also be used to migrate to Exadata. Any technique used to move data between Oracle Databases can be used with Exadata.

NO.4 Which two statements are true about the use ofi direct path loads when selecting firom external tables in a database
on a Database Machine?
A. INSERT INTO . . . SELECT FROM statements, executed serially, which select firom external tables, require the
APPEND hint to use direct path loading.
B. CREATE TABLE . . . AS SELECT statements, which select firom external tables, attempt to use in direct path loading
automatically.
C. CREATE TABLE . . . AS SELECT statements, which select firom external tables, require the APPEND hint to use direct
path loading.
D. INSERT INTO . . . SELECT FROM statements, executed serially, which select firom external tables, are unable to use
direct path loading.
Answer: A, B
Explanation:
A CTAS (Create table as select) will always use direct path (B, not C) load but IAS (Insert as select) statement will not.
In order to achieve direct path load with an IAS statement you must add the APPEND hint to the command (A, not D).
Direct path loads can also run in parallel. You can set the parallel degree fior a direct path load either by adding the PARALLEL hint to the CTAS or IAS statement or by setting the PARALLEL clause on both the external table and the table into which the data will be loaded. Once the parallel degree has been set at CTAS will automatically do direct path load in parallel but an IAS will not. In order to enable an IAS to do direct path load in parallel you must alter the session to enable parallel DML.
Note:
* Parallel Direct Path Load
The key to good load perfiormance is to use direct path loads wherever possible. A direct path load parses the input data according to the description given in the external table definition, converts the data fior each input field to its corresponding Oracle data type, then builds a column array structure fior the data. These column array structures are used to fiormat Oracle data blocks

and build index keys. The newly fiormatted database blocks are then written directly to the database, bypassing the standard SQL processing engine and the database buffer cache.
Refierence: Best Practices fior Implementing a Data Warehouse on the Oracle Exadata Database Machine; Using CTAS &
Exchange Partition Replace IAS fior Copying Partition on Exadata

NO.5 Which two are true about Smart Scan?
A. a query rewrite may occur to a container table stored in Exadata but will never benefit From Smart scan.
B. Column projection does not contribute to the perfiormance benefit ofi Smart Scan
C. It is possible to offload single row fiunctions to the storage servers.
D. Some joins can be offloaded to the storage servers.
E. A query rewrite may occur to a container table stored Exadata, and it will always benefit firom Smart Scan.
F. All joins can be offloaded to the storage servers.
Answer: C, D
Explanation:
C: With Exadata storage, database operations are handled much more efficiently. Queries that perfiorm table scans
can be processed within Exadata storage with only the required subset ofi data returned to the database server. Row
filtering, column filtering and some join processing (among other fiunctions) are perfiormed within the Exadata storage
cells. When this takes place only the relevant and required data is returned to the database server.
D (not F):
* Exadata perfiorms joins between large tables and small lookup tables, a very common scenario fior data warehouses
with star schemas. Joining large tables and small lookup tables is implemented using Bloom Filters, which are a very
efficient probabilistic method to determine whether a row is a member ofi the desired result set.
* Ifi storage indexes are so great, why doesn't Oracle Exadata use them all the time? The short answer is that they are
created and used only when they will be beneficial.
* To use storage indexes, Oracle Exadata queries must use smart scans, so not all types ofi applications can benefit
firom storage indexes. Applications with queries that include predicates and perfiorm a lot ofi fiull table scans or fiast fiull
scans ofi indexes-typically those used in data warehousing environments-will benefit greatly firom storage indexes.
Online transaction processing (OLTP) applications, on the other hand, typically access a small number ofi rows through
standard indexes and do not perfiorm fiull table scans, so they may not benefit firom storage indexes
.
Note:
* Storage indexes reside in the memory ofi the storage servers-also called storage cells-and significantly reduce

unnecessary I/O by excluding irrelevant database blocks in the storage cells.
* To use storage indexes, Oracle Exadata queries must use smart scans, so not all types ofi applications can benefit
firom storage indexes.
Incorrect:
Not B: Exadata provides column filtering, also called column projection, fior table scans. Only the columns requested
are returned to the database server rather than all columns in a table. For example, when the fiollowing SQL is issued,
only the employee_name and employee_number columns are returned firom Exadata to the database kernel.
SELECT employee_name, employee_number FROM employee_table.
For tables with many columns, or columns containing LOBs (Large Objects), the I/O bandwidth saved can be very
large. Using both predicate and column filtering dramatically improves perfiormance and reduces I/O bandwidth
consumption. In addition, column filtering also applies to indexes, allowing fior even fiaster query perfiormance.
Refierence: Oracle Communications Data Model Implementation and Operations Guide, Exadata Smart Scan
Processing and Storage Index

NO.6 Which two are true about the use ofi DBFS in a Database Machine environment?
A. DBFS must be used to bulk load data into a database on the Database Machine ifi the staging area requires Exadata
based shared storage.
B. DBFS must be used to have a POSIX compliant shared storage solution that is accessible firom the database servers
on a Database Machine.
C. DBFS must be used to bulk load data into a production database on the Database Machine.
D. DBFS must use the DBFS_DG diskgroup fior any DBFS store.
E. DBFS must be used to have a POSIX-compliant Exadata-based shared storage solution.
Answer: CD