

## 1Z0-027 Dumps

# Oracle Exadata Database Machine Administration, Software Release 11.x

<https://www.certleader.com/1Z0-027-dumps.html>



**NEW QUESTION 1**

Which tool will provide you with diagnostic information for all the software log, trace files, and OS information on Database Machine?

- A. dbmcheck.sh
- B. diagget.sh
- C. oswatcher
- D. adrci
- E. Enterprise Manager

**Answer: B**

**Explanation:**

Gather all diagnostics information  
/opt/oracle.SupportTools/onecommand/diagget.sh

**NEW QUESTION 2**

Identify two valid reasons for creating multiple griddisks on a Single celldisk.

- A. To segregate storage into multiple pools with different performance characteristics
- B. To facilitate normal or high redundancy ASM diskgroups
- C. To enable disk mirroring for the system area
- D. To segregate storage into multiple pools that can be assigned to different databases
- E. To segregate storage into multiple pools that can be assigned to different resource consumer groups in the same database.

**Answer: AD**

**NEW QUESTION 3**

You plan to migrate an existing production database supporting online transaction processing (OLTP) workloads to the Exadata Database Machine. The database currently supports an application requiring fast response times, which satisfies stringent requirements, and most of the application queries use index access to the tables in the application schema.

For which case would you consider dropping indexes to allow Smart Scans to occur?

- A. Drop non-constraint indexes if Smart Scan occurs instead of an index access path on the corresponding table.
- B. Drop non-constraint indexes if Smart Scan performs better than index unique scans on the corresponding table.
- C. Drop non-constraint indexes if Smart Scan performs better than index range scans on the corresponding table.
- D. Drop non-constraint Indexes if Smart Scan performs better than any index scans on the corresponding table.

**Answer: D**

**NEW QUESTION 4**

You are examining the existing IORM configuration on the cells of Database Machine, to see if they require my modifications based on recent changes to various workloads.

All seven cells in your X3-2 half-rack shown the following:

```
CellCLI> list iormplan detail
name:          dmor1cel04_IORMPLAN
catPlan:       name=interactive,level=1,allocation=90
               name=batch,level=2,allocation=80
               name=maintenance,level=3,allocation=50
               name=other,level=3,allocation=50
dbPlan:        name=sales,level=1,allocation=45,flashcache=on,flashlog=on
               name=finance,level=1,allocation=45,flashcache=on,flashlog=off
               name=other,level=1,allocation=10,flashcache=off,flashlog=off
objective:     off
status:        active
```

Which two are true about I/O to the cells using this plan?

- A. I/O requests in the batch category may use flashcache if the I/O is from the sales finance database, and these I/O requests are guaranteed to get 80% of the I/O if the interactive category I/Os use no more than 20%.
- B. I/O requests made by sessions in the marketing database may use flashing and flashcache if no other categories or database or database are using flashing and flashcache at the same time.
- C. I/O requested in the interactive category may use flashdns if the I/O is from the sales or finance databases, and these I/O requests are guaranteed to get 90% of the I/O if the enough I/Os are issued in this category.
- D. I/O requests from the sales database may use flashing regardless of the I/O category.
- E. No I/Os in any category or from any database may use flashing or flashcache because the objective is off.

**Answer: AD**

**NEW QUESTION 5**

You plan to migrate your Oracle Version 11.1.0.2 database to your Exadata Database Machine.

The database supports an online transaction processing (OLTP) workload and is currently hosted on a Little Endian platform

Which two are the supported and appropriate migration methods to minimize downtime?

- A. Upgrade source database to 11.2.0 and migrate using a physical standby database.
- B. Migrate using Data Pump.
- C. Migrate using GoldenGate.
- D. Migrate using cross platform Transportable Database.
- E. Migrate using ASM online migration.

**Answer: AB**

**Explanation:**

Reference: Migrating the Oracle E-Business Suite Database to Oracle Exadata Database Machine Using Transportable Tablespaces  
Reference: Migrating Oracle E-Business Suite to Oracle Exadata Database Machine Using Oracle Data Pump

**NEW QUESTION 6**

Identify the three components that serve a purpose only in the Database Machine.

- A. ASM intelligent Data Placement (IDP)
- B. Intelligent Database Protocol (IDB)
- C. Database Resource Manager (DBRM)
- D. I/O Resource Manager (IORM)
- E. Database Filesystem (DBFS)
- F. The DISKMON process

**Answer:** ABD

**Explanation:**

Intelligent Data Placement, a feature of ASM that allows placing data in such a way that more frequently accessed data is located close to the periphery of the disk where the access is faster.

The Exadata software is optimally divided between the database servers and Exadata cells. The database servers and Exadata Storage Server Software communicate using the iDB –

the Intelligent Database protocol. iDB is implemented in the database kernel and transparently maps database operations to Exadata-enhanced operations. iDB implements a function shipping architecture in addition to the traditional data block shipping provided by the database. iDB is used to ship SQL operations down to the Exadata cells for execution and to return query result sets to the database kernel. Instead of returning database blocks, Exadata cells return only the

The inter-database I/O allocations are defined within the software in the Exadata cell and managed by the I/O Resource Manager (IORM). The Exadata cell software ensures that inter-database I/O resources are managed and properly allocated within, and between, databases.

**NEW QUESTION 7**

Your database Machine has the exachk utility pre-installed and you decide to use it periodically, to validate the installation against Oracle's recommended best practices.

Which two actions could you take to do this?

- A. Use a cron job on a database node to run it at regular intervals.
- B. Run it once from a database node and it will then perform periodic monitoring automatically.
- C. Use a cron job on each cell to run it at regular intervals.
- D. Run it once on each cell and it will then perform periodic monitoring automatically.
- E. Create a Job in Enterprise Manager to run the exachk utility at regular intervals.

**Answer:** AE

**NEW QUESTION 8**

You must apply patches and patch bundles in a rolling fashion, if possible, on the components of your Database Machine.

You use RAC for your database and also use Data Guard, having standby database on another Database Machine.

You wish to have scripts that contain the appropriate command 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 the following:

```
$ORACLE_HOME/oplan/oplangenerateApplySteps /u01/stage
```

Which two are true concerning oplan?

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

**Answer:** CE

**Explanation:**

Oplan generates instructions for all of the nodes in the cluster.

Note:

\* Oplan is a utility that facilitates you with the application of bundle patches on Exadata compute nodes via Opatch. This new utility helps you with the patch process by generating step-by-step instructions telling you how to apply a bundle patch in your environment.

\* Exadata is much more than a "database-in-a-box." It is a set of compute nodes (think RAC node servers) combined with ultra-fast infiniband (...and 10GB ethernet, and multiple

\* 1GB ethernet) and storage nodes.

Incorrect:

Not D: There is no support for Oracle DataGuard

**NEW QUESTION 9**

You are in the planning stage of the network configuration for your Database Machine. The requirements are:

1. A fault-tolerant network, providing higher availability for connections to database instances
2. Fault tolerance providing higher availability for connections to perform management functions on the database and storage servers.
3. Full monitoring of all Database Machine components using Enterprise Manager

Which three components require external Ethernet network cables to connect your existing network infrastructure to your database machine to satisfy this requirement?

- A. Database servers
- B. Exadata storage servers

- C. InfiniBand storage servers.
- D. Power distribution units (PDUs)
- E. Cisco Ethernet switch

**Answer:** ADE

**NEW QUESTION 10**

Which statement is true about operating systems on database Machine multi-rack configuration consisting of two full racks and one Exadata storage rack?

- A. All Exadata storage servers and all database servers must run the Oracle Linux O/S and X3-8 database server may run the Oracle Solaris O/S.
- B. All Exadata storage servers must run the Oracle Linux O/S and all database servers within the same cluster must run the Oracle Linux O/s.
- C. All Exadata storage servers must run the Oracle Linux O/S and all database servers within the same cluster must run the same O/s.
- D. All Exadata Storage Servers must run the Oracle Solaris O/S and all database servers within the same cluster must run the same O/S.
- E. All Exadata storage servers in the same cluster must run the same O/S but Exadata Storage Servers in different clusters may run a different O/S.

**Answer:** C

**NEW QUESTION 10**

Which two are true about the use of DBFS in a Database Machine environment?

- A. DBFS must be used to bulk load data into a database on the Database Machine if the staging area requires Exadata based shared storage.
- B. DBFS must be used to have a POSIX compliant shared storage solution that is accessible from 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 for any DBFS store.
- E. DBFS must be used to have a POSIX-compliant Exadata-based shared storage solution.

**Answer:** CD

**NEW QUESTION 12**

Which three storage components are available after the standard initial Database machine deployment?

- A. The DATA\_<DBM\_Name> ASM diskgroup
- B. The RECO\_<DBM\_Name> ASM diskgroup
- C. Mirrored system partitions on hard disk 0 and hard disk 1
- D. The DBFS\_DG diskgroup with external redundancy
- E. Exadata Smart Flash Cache using all of the flashdisk space

**Answer:** ABC

**NEW QUESTION 17**

You are about to run the oplan utility to patch the servers on your test Database Machine before patching the production environment.

The following task might be performed:

- A) Test the failback procedure
- B) Run the exachk utility
- C) Read the README file.
- D) Automate the patch application process as appropriate.
- E) Verify that the patch provides the functionality it is meant to.
- F) Apply the patch.
- G) Evaluate the system performance.

In which order should the tasks be performed to patch in the recommended fashion?

- A. C, B, D, F, B, E, A
- B. C, D, F, B, E, G, A, B
- C. C, B, D, F, E, G, A
- D. C, B, D, F, E, A, G
- E. C, B, D, F, B, E, G, A

**Answer:** E

**NEW QUESTION 21**

Your Database Machine has a large database with some very large tables supporting OLTP workloads.

High volume Insert applications and high volume update workloads access the same tables.

You decide to compress these tables without causing unacceptable performance overheads to the OLTP application.

Which three are true regarding this requirement?

- A. Using 'compress for oltp' will compress the data less than if using Hybrid Columnar Compression when specified with compress for query low.
- B. The compression is performed on the storage servers when using compress for oltp in an Exadata environment.
- C. The compression method compress for archive high is the worst fit for this requirement.
- D. Using 'compress for oltp' will compress the data more than if using Hybrid Columnar Compression when specified with compress for archive low.
- E. The compression is performed on the database servers when using 'compress for oltp' in an Exadata environment.

**Answer:** ACE

**Explanation:**

Note:

(E not B):

\* Types of compression

Basic compression OLTP compression Warehouse compression

Online archival compressio

\*

/ OLTP compression allows compression during DML operations.

/ Basic compression works at the data block level.

\* When you enable table compression by specifying COMPRESS FOR OLTP, you enable OLTP table compression. Oracle Database compresses data during all DML operations on the table. This form of compression is recommended for OLTP environments.

\* When you specify COMPRESS FOR QUERY or COMPRESS FOR ARCHIVE, you enable hybrid columnar compression. With hybrid columnar compression, data can be compressed during bulk load operations. During the load process, data is transformed into a column-oriented format and then compressed. Oracle Database uses a compression algorithm appropriate for the level you specify. In general, the higher the level, the greater the compression ratio. Hybrid columnar compression can result in higher compression ratios, at a greater CPU cost. Therefore, this form of compression is recommended for data that is not frequently updated.

### NEW QUESTION 22

Which two activities are supported on the storage servers in the Database Machine?

- A. Installing an alternative package manager
- B. configuring secure shell user equivalency for the callmonitor user
- C. changing root password
- D. upgrading the Storage Server software package using RPN
- E. upgrading a device driver for hard disks when inserting a replacement disk after a hard disk failure

**Answer: BC**

### NEW QUESTION 24

You are monitoring and evaluating a create index statement on your Database Machine and have run the following query after executing statement, providing the output. Shown:

```
SQL> SELECT DISTINCT event, total_waits, time_waited/100 wait_secs,
2 average_wait/100 avg_wait_secs
3 FROM V$SESSION_EVENT e, V$MYSTAT s
4 WHERE event LIKE 'cell%' AND e.sid = s.sid;
```

EVENT	TOTAL_WAITS	WAIT_SECS	AVG_WAIT_SECS
cell list of blocks physical read	1	0	.0006
cell single block physical read	1349704	683.94	.0005
cell smart table scan	9191	3.29	.0004

Select two reasons why the statement would have produced so many “cell single block physical read” waits compared to “cell smart table scan” waits.

- A. There are huge numbers of migrated rows in the table on which the index is being built.
- B. There is an uncommitted transaction that has modified one block of the table on which the index is being built, in each cell.
- C. There is a transaction that has modified one block of the table on which the index is being built in each cell, which committed after the create index began.
- D. There are huge numbers of chained rows in the table on which the index is being built.
- E. There is a ROWID column in the table on which the index is being built.

**Answer: AD**

#### Explanation:

A: It could be that row migration.

D: It could be that row migration or chained rows could cause it.

Note:

\* Some facts about scans: Scans exists in “OLTP” systems

Exadata smart scan requires a direct path read.

A direct path read is chosen at runtime based on internal heuristics The STORAGE clause in an explain plan doesnt necessarily mean you will perform a smart scan.

\* The buffer caching in certain “OLTP” environments can occasionally induce conventional reads when smart scan is faster.

\* Typically see cell multiblock physical read instead of cell smart table scan waits

\* No one-size-fits-all solution can be given here but it is very fixable.

### NEW QUESTION 25

Which two communication methods are used by which components in the Enterprise manager Architecture for the Database Machine?

- A. SNMP traps for alerts are sent by the storage server ILOM to the storage server MS process
- B. SNMP traps for alerts are sent by the storage server MS process to the storage server ILOM
- C. SNMP traps for alerts are sent by the storage server ILOM to the Enterprise Manager agent.
- D. SNMP traps for alerts are sent by the storage server MS process to the enterprise Manager agent
- E. SNMP traps for alerts are sent by the storage server ILOM to the storage server RS process.

**Answer: AD**

#### Explanation:

There are two types of server alerts that come from Oracle Exadata Storage Server:

\* (A) For Integrated Lights Out Manager (ILOM)-monitored hardware components, ILOM reports a failure or threshold exceeded condition as an SNMP trap, which is received by MS.

MS processes the trap, creates an alert for the storage server, and delivers the alert via SNMP to Oracle Enterprise Manager 12c.

\* (D) For MS-monitored hardware and software components, MS processes a failure or threshold exceeded condition for these components, creates an alert, and delivers the alert via SNMP to Oracle Enterprise Manager Cloud Control 12c.

Reference: Managing Oracle Exadata with Oracle Enterprise Manager 12c, Oracle White Paper

### NEW QUESTION 26

Which two statements are true about the IPTables firewall configuration on a Database Machine- Machine after the default Initial deployment?

- A. IPTables is configured with Oracle supplied rules on the cells.
- B. IPTables is configured with Oracle supplied rules on the database servers.
- C. IPTables is installed and available but not configured on any servers.
- D. IPTables is installed and available but not configured on the database servers.
- E. IPTables is installed and available but not configured on the cells.

**Answer:** AD

**NEW QUESTION 30**

Which two are true concerning the allocation of I/O resources by the IORM within the CELLSRV process?

- A. Control File I/O is managed automatically at high priority by IORM.
- B. Control File I/O is considered part of the SYSTEM resource Consume group by IORM.
- C. Log Writer I/O to the Smart Flash Log is considered part of the SYSTEM resource consumer group by IORM.
- D. Log Writer I/O to the Smart Flash Log is managed automatically at high priority by IORM.
- E. Database Writer I/O is managed automatically at normal priority by IORM.
- F. Database Writer I/O is considered part of the SYSTEM resource consumer group by IORM.

**Answer:** AE

**Explanation:**

Note:

\* IORM Rules

IORM is only "engaged" when needed.

/ (A) Redo and control file writes always take precedence.

/ (E) DBWR (database writer) writes are scheduled at the same priority as user IO.

/ IORM does not intervene if there is only one active consumer group on one database.

/ Any disk allocation that is not fully utilized is made available to other workloads in relation to the configured resource plans.

/ Background IO is scheduled based on their priority relative to user IO.

/ For each cell disk, each database accessing the cell has one IO queue per consumer group and three background queues.

/ Background IO queues are mapped to "high", "medium", and "low" priority requests with different IO types mapped to each queue.

/ If no intradatabase plan is set, all non-background IO requests are grouped into a single consumer group called OTHER\_GROUPS.

Reference: Using IORM with Exadata

**NEW QUESTION 31**

.....

## Thank You for Trying Our Product

\* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

\* One year free update

You can enjoy free update one year. 24x7 online support.

\* Trusted by Millions

We currently serve more than 30,000,000 customers.

\* Shop Securely

All transactions are protected by VeriSign!

**100% Pass Your 1Z0-027 Exam with Our Prep Materials Via below:**

<https://www.certleader.com/1Z0-027-dumps.html>