

## 1Z0-027 Dumps

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

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



#### NEW QUESTION 1

You have used setupem.sh to deploy a plug in for Grid Control.  
In which two ways do all the targets supported by the plug-in get configured?

- A. The targets must be configured by an Enterprise Manager administrator manually using Grid Control.
- B. Setupem.sh is used again to configure the targets.
- C. The targets must be configured by an O/S administrator manually.
- D. The targets must be configured by a database administrator manually.
- E. There may be more than one target for each plug-in.

**Answer:** AE

#### NEW QUESTION 2

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 3

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

- A. If two consumer groups A and B in the PROD database get 10% and 15% respectively of resource allocation, then the percentage of I/O resource would be the same if they got 20% and 30% respectively, provided that the category plans and inter database plans are unchanged, and consumer groups A and B are still mapped to the same category.
- B. If two consumer groups C and D in the PROD database get 10% and 15% respectively of resource allocation, then the percentage of I/O resource would be the same if they still got 10% and 15% respectively and were remapped to a different category by the DBA, provided that the category plans and Interdatabase plans are unchanged.
- C. If two consumer groups P and Q in the PROD database get 20% and 10% respectively of resource allocation, then the percentage of I/O resource would be the same if they got 10% and 5% respectively, and the Interdatabase plan has changed, provided that the category plan is unchanged, and consumer groups A and B are still mapped to the same categories.
- D. Log Writer I/O and Control File I/O are high priority I/O managed automatically by IORM.
- E. Database Writer I/O is considered to be the highest priority I/O and is managed based on attributes of the IORM plan.

**Answer:** CD

#### Explanation:

\*(not A, Not B) Rules in an interdatabase resource plan specify allocations to databases, not consumer groups.  
\* IORM Rules  
IORM is only "engaged" when needed.  
/ (D) Redo and control file writes always take precedence.  
/ (Not 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  
Note:  
\* Use IORM metrics to track  
/ I/O load per Consumer Group (IOPS, MBPS, disk utilization %)  
/ I/O throttling per Consumer Group Reference: Using IORM with Exadata

#### NEW QUESTION 4

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 5

A read-only application is in development and is using a test database on a Database Machine. You are examining SQL statements from this application in an attempt to determine which ones will benefit from the Exadata Smart scan capability.

The following is true about the tables used by the application:

1. The data for the tables has just been loaded.
2. There are no applications accessing the tables currently.

3. None of the indexes are compressed or reverse key indexes.
  4. The tables use the default organization type data.
  5. The only data types for the table are varchar (2), number, or date.
  6. The largest number of columns for any table is 29.
  7. No access is based on ROWID, or virtual columns.
- Which two access paths will always generate either a set of "cell smart table scan" or a set of "cell smart index scan" requested?

- A. Full scans on sorted hash clustered tables executed in parallel
- B. Full table scans on index organized tables executed in parallel.
- C. Full table scans on heap tables executed in parallel
- D. Full scans on index clustered tables executed in parallel
- E. full scans on hash clustered tables executed in serially
- F. fast full scans on B\*Tree indexes executed in parallel
- G. full index scans on B\*Tree indexes executed in parallel

**Answer:** CF

#### NEW QUESTION 6

You have two very large databases supporting OLTP workloads which run on multiple small-blade style servers in a cluster and which require optimal latency for I/O.

You plan to migrate them to a Database Machine once a capacity planning exercise is finished.

The database backup strategy requires that backups are written directly to media.

High availability requirements state that you must be able to survive node failures at any time.

Which three Database Machine components or features would you recommend to support these requirements?

- A. Use of write back flash cache
- B. Use of smart flash logs
- C. High capacity disks in the Database Machine
- D. High performance disks in the Database Machine
- E. A high capacity expansion full rack
- F. A high performance expansion full rack

**Answer:** ABC

#### NEW QUESTION 7

Consider the following list of software components:

1. DCLI
2. Management Server (MS)
3. ASM Instance
4. RDBMS instance
5. Restart Server (RS)
6. Cellcli
7. Cell Server (CELLSRV)
8. Diskmon

Identify the location where these software components may run in the standard Database machine deployment.

- A. 3, 4 and 8 run on the database servers; 1, 2, 5, 6 and 9 run on the Exadata Storage servers.
- B. 4 and 8 run on the database servers; 1, 2, 3, 5, 6 and 7 run on the Exadata Storage servers.
- C. 1, 3 and 4 run on the database servers; 2, 5, 6, 7 and 8 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. 3, 4 and 8 run on the database servers; 1, 2, 5, 7 and 8 run on the Exadata Storage servers.
- F. 1, 3, 4 and 8 run on the database servers; 1, 2, 5, 7 and 8 run on the Exadata Storage servers.

**Answer:** D

#### NEW QUESTION 8

Which four statements are true about Exadata Smart Flash Cache?

- A. Smart Scan will always be done for I/Os to flash based griddisks.
- B. Flash based ASM diskgroups may share space with the Flash Cache on the flashdisks.
- C. Single block reads can benefit from Smart Flash Cache.
- D. Smart Scan will never be done for I/Os to flash based griddisks.
- E. Multiblock reads can benefit from Smart Flash Cache.
- F. Smart Flash Logs reduce the size of Smart Flash Cache.

**Answer:** BCEF

#### Explanation:

B: \* Grid disks (the logical disks

that reside on physical cell disks) are created on these flash-based cell disks and the grid disks are assigned to an Automatic Storage Management (ASM) diskgroup. The best practice would be to reserve the same amount of flash on each Exadata cell for flash disks and have the ASM diskgroup spread evenly across the Exadata cells in the configuration just as you would do for regular Exadata grid disks. This will evenly distribute the flash I/O load across the Exadata cells and flash.

Note:

\* The Exadata

Storage Server comes with a substantial amount of flash storage. A small amount is allocated for database logging and the remainder will be used for caching user data

#### 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 three statements are true regarding the configuration of Auto Service Request (ASR) on your Database Machine?

- A. HTTPS connectivity must be enabled from the ASK Manager host to the internet.
- B. ASR Manager must be installed on one of the database servers.
- C. Oracle Solaris must be used on the server where ASK Manager is installed.
- D. SMTP must be enabled on at least one database server.
- E. Oracle Linux must be used on the server where ASR Manager is installed
- F. ASR Manager may be installed on any type of server running Oracle Linux.
- G. ASR Manager may be installed on any type of server running Oracle Solaris.

**Answer:** AFG

#### Explanation:

A: Before installing ASR, please ensure the following conditions are met:

/ (A) Ensure connectivity to the Internet using HTTPS.

/ Make sure you have access to My Oracle Support and that your contact information is correct and current.

/ Make sure all of your assets have a Contact assigned and that the contact is correct and current.

/ Identify and designate a system to serve as ASR Manager.

/ Identify and verify ASR assets.

F, G (not C, not E, not B): The recommended configuration is to install the ASR Manager, which receives fault telemetry information from the servers in Oracle Exadata Database Machine, on an external standalone server. This server must run Solaris or Linux as the operating system.

Reference: Oracle Auto Service Request Exadata Database Machine Quick Installation Guide

#### NEW QUESTION 11

Which is true regarding Enterprise Manager monitoring and configuration?

- A. Agents must be deployed on the- storage servers, in case all the database servers are down.
- B. Monitoring the power distribution units (PDUs) is done by Enterprise Manager agents, but additional ethernet cables are required to connect the PDUs to the internal network switch in the Database Machine
- C. Monitoring the power distribution units (PDUs) is done by Enterprise Manager agents without, any additional network configuration.
- D. Agents may be deployed on the storage servers, in case all the database servers are down.
- E. Monitoring the power distribution units (PDUS) is done by Enterprise Manager agents but additional ethernet cables are required to connect the PDUs to the corporate network switch.

**Answer:** C

#### Explanation:

Note:

\* The following prerequisites must be met before you can deploy the plug-in:

\* Set and validate the preferred credentials on all Agents where you want to deploy the plug- in. This is essential in order to deploy the plug-in.

The PDU firmware version must be 1.02 or later.

\* This plug-in will be used to monitor PDU's actual current value of Phase1, Phase2, and Phase3 and raises appropriate events if they have crossed the different threshold values. Each PDU will have a plug-in instance added to its own OEM Grid Control agent.

By deploying the plug-in to your Grid Control environment, you gain the following management features:

Monitor PDU targets.

Raise alerts and violations based on thresholds set on monitoring and configuration data. Provide rich out-of-box metrics and reports based on the gather

Reference: Oracle Enterprise Manager System Monitoring Plug-In Installation Guide for Exadata Power Distribution Unit

#### NEW QUESTION 12

Last weekend, an Exadata storage server flashdisk entered the predictive failure state. The flashdisk is used by the flashcache and has a griddisk which is a member of a normal redundancy diskgroup. Identify the four steps you must perform to replace this flashdisk.

- A. Identify the griddisk on the predictive failure flashdisk and drop it from the associated ASM diskgroup
- B. Verify that the griddisk located on the predictive failure flashdisk has been successfully dropped from the associated ASM diskgroup.
- C. Drop the flashcache on the cell and re-create it using all but the predictive failure flashdisk.
- D. Safely power off the cell containing the predictive failure flashdisk.
- E. Replace the predictive failure flashdisk.
- F. Power up the cell containing the replaced flashdisk and activate all griddisks.
- G. Drop the flashcache on the cell and re-create it using all flashdisks.
- H. Create a new griddisk on the replaced flashdisk.
- I. Add the griddisk back into the ASM diskgroup to which it belonged.

**Answer:** ADEF

#### NEW QUESTION 16

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 17

Which two may act as building blocks for the creation of two or more griddisks residing on the same physical device and which have more balanced performance characteristics?

- A. griddisk on hard drive-based interleaved celldisks.
- B. griddisk on hard drive-based celldisks consumed by ASM when creating diskgroups using ASM-based intelligent Data Placement (IDP)
- C. griddisk on hard drive-based non-interleaved celldisks.
- D. griddisk on Flash-base celldisks.

**Answer:** AD

#### Explanation:

Interleaved Griddisks mean equally fast Griddisks, whereas with non Interleaved Griddisks (the default) the Griddisks created first are being taken from the outer tracks of the underlying physical disk.

#### NEW QUESTION 22

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 26

Which three are true regarding the use of Storage Indexes?

- A. Different storage regions may have different columns indexed for the same table.
- B. A Storage index is automatically maintained by CELLSRV based on the filter columns of the offload SQL.
- C. The use of Storage indexes for a particular database can be disabled by using an I/O Resource Manager Database Plan.

- D. Storage Indexes occupy space in the Smart Flash Cache.
- E. The use of Storage Indexes for particular categories of I/O can be disabled by using an I/O Resource Manager Category Plan.
- F. A maximum of eight table columns for any table are Indexed per storage region.

**Answer:** ABF

**Explanation:**

Note:

- \* Storage indexes are used during smart scans. All the limitations to smart scans apply to storage indexes. They do not work with joins. Bind variables are supported, however it's slightly more restrictive than regular indexes/queries.
- \* The storage index is stored in the memory on each of the Exadata storage cells and is created and maintained transparently. However, if a storage cell is shutdown or rebooted the storage index will be lost from memory and will be recreated on subsequent accesses to the data after the cell has been brought back online.
- \* Storage Indexes are a very powerful capability provided in Exadata storage that helps avoid I/O operations. The Exadata Storage Server Software creates and maintains a Storage Index (that is, metadata about the database objects) in the Exadata cell. The Storage Index keeps track of minimum and maximum values of columns for tables stored on that cell. When a query specifies a WHERE clause, but before any I/O is done, the Exadata software examines the Storage Index to determine if rows with the specified column value exist in the cell by comparing the column value to the minimum and maximum values maintained in the Storage Index. If the column value is outside the minimum and maximum range, scan I/O for that query is avoided. Many SQL Operations run dramatically faster because large numbers of I/O operations are automatically replaced by a few lookups. To minimize operational overhead, Storage Indexes are created and maintained transparently and automatically by the Exadata Storage Server Software.

**NEW QUESTION 29**

You are using Hybrid Columnar Compression for a table stored in a tablespace that is contained in an Exadata-based ASM diskgroup. Identify three statements that correctly explain where the compression and decompression can be done.

- A. Decompression can be done on the database servers.
- B. Compression can be done on the Exadata storage servers.
- C. Compression can be done on the database servers.
- D. Decompression can be done on the Exadata storage servers.

**Answer:** ACD

**Explanation:**

- \* decompression
- / Queries run directly on Hybrid Columnar Compressed data does not require the data to be decompressed
- / Data that is required to satisfy a query predicate does not need to be decompressed; only the columns and rows being returned to the client are decompressed in memory
- / The decompression process typically takes place on the Oracle Exadata Storage Server in order to maximize performance and offload processing from the database server.

**NEW QUESTION 33**

Which two are true about the use of the Integrated Lights Out Manager (ILOM) on the Database Machine?

- A. ILOM can be used to power-on the Cisco switch.
- B. ILOM can be used to power on the InfiniBand switches.
- C. ILOM can be used to power-on the database servers.
- D. ILOM generates hardware alerts for the power distribution units.
- E. ILOM provides a remote console for the storage servers.

**Answer:** CE

**Explanation:**

- \* The Oracle ILOM Remote Console is supported on all Oracle Sun x86 processor-based servers. It is also supported on some SPARC processor-based servers. The Oracle ILOM Remote Console is a Java application that you can launch from the Oracle ILOM web interface
- \* Remote Management: Integrated Lights Out Manager (ILOM) Ethernet port

**NEW QUESTION 37**

You plan to monitor storage servers after configuring an I/O resource manager plan with directives for inter-database plans and intra-database plans. Which two types of metrics would help assess the impact of the intra-database plans on I/O to the storage servers?

- A. Category I/O
- B. Database I/O
- C. Resource Consumer Group I/O
- D. Smart Flash Log I/O
- E. Smart Flash Cache I/O

**Answer:** BC

**Explanation:**

- B: Database metrics provide information about the size of the I/O load from each database specified in the interdatabase plan.
  - C: Consumer group metrics provide information about the size of the I/O load from each consumer group specified in a database resource plan. Each database in the interdatabase plan has metrics for each of its consumer groups.
- Note:
- \* I/O Resource Manager (IORM) Settings
  - \* Incorrect:
- Not A: Category metrics provide information about the size of the I/O load from each category specified in the current IORM category plan.

**NEW QUESTION 40**

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 41**

You must drop all cell disks on all the cells in a quarter rack as part of a reconfiguration project, to support normal redundancy interleaving. Select two actions that describe the operating system (O/S) account on the cells to which you should log in, and the tool that may be used to drop the cell disks.

- A. To the CELLMONITOR account using CELLCLI interactively on each cell.
- B. To CELLADMIN account calling CELLCLI on all cells using DCLI
- C. To either the CELIMONITOR or the CELLADMIN account calling CELLCLI on all cells using DCLI
- D. To the CELLMONITOR account calling CELLCLI on all cells using DCLI
- E. To the CELLADMIN account using CELLCLI interactively on each cell
- F. To either the CELLMONITOR or CELLADMIN account calling interactively on each cell

**Answer:** BE

**Explanation:**

Note:

\* Each Exadata Storage cell can be monitored, configured, and maintained using the cellcli command line interface. To invoke the CLI, simply login to the Exadata cell as cellmonitor, celladmin, or root, and type "cellcli". Within cellcli, a wide range of commands can be invoked to monitor the cell.

\* While connected to cellcli as the Linux account celladmin, you can perform all cellcli commands except calibrate.

While logged into the cell server as cellmonitor, you can do any "read-only" commands; you cannot change any configuration or attributes.

**NEW QUESTION 46**

.....

## 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>