2023 1Z0-902 dumps review - Professional Quiz Study Materials [Q25-Q47

2023 1Z0-902 dumps review - Professional Quiz Study Materials
1Z0-902 Test Prep Training Practice Exam Questions Practice Tests

NO.25 Which three of the following options are available for database consolidation on Exadata?
* Multiple pluggable databases in one or more container databases on a Virtual Machine or Bare Metal cluster
* Bare Metal deployment with one database per database server
* A single database with one or more schema per application
* Multiple databases spanning a Bare Metal and Virtual Machine cluster to provide resource isolation and consolidation concurrently
* Multiple Databases on a cluster
* Multiple Virtual Machine clusters supporting a single database for greatest resource isolation

According to Oracle.com documents or resources, the three options that are available for database consolidation on Exadata are:

A) Multiple pluggable databases in one or more container databases on a Virtual Machine or Bare Metal cluster123.

C) A single database with one or more schema per application12.

E) Multiple Databases on a cluster12.

https://blogs.oracle.com/exadata/post/database-consolidation-why-and-how

NO.26 What are two recommended configuration best practices for backup and recovery on Exadata?
* Placing the backup network on dedicated switches installed in the top of the rack (ToR) has the benefits of isolating the backup network from other workloads and providing a greater level of control.
* Even if the backup was limited to writing to four tape drives, eight channels could be specified to expedite the restore process.
* The internal Recovery Appliance backup and restore processing is optimized when the RMAN FILESPERSET parameter is set to 1 for the level 1 incremental backup set.
* When off-site long-term backup retention is needed, use Oracle Database Backup Cloud Service as a low-cost, offsite scalable storage for a disaster recovery solution.
According to the Oracle Exadata Database Machine documentation12, two recommended configuration best practices for backup and recovery on Exadata are:

1. Placing the backup network on dedicated switches installed in the top of the rack (ToR) has the benefits of isolating the backup network from other workloads and providing a greater level of control.

2. When off-site long-term backup retention is needed, use Oracle Database Backup Cloud Service as a low-cost, offsite scalable storage for a disaster recovery solution.

NO.27 Which two activities are supported on the storage servers in an Exadata Database Machine?
* changing the root password
* upgrading a device driver for hard disks when inserting a replacement disk after a hard disk failure
* installing an alternative package manager
* upgrading the Storage Server software package using rpm
* configuring secure shell user equivalency for the cellmonitor user
According to the web search results123, the storage servers in an Exadata Database Machine are mainly used for processing data at the storage level and offloading some SQL operations from the database servers. Therefore, the two activities that are supported on the storage servers are:

A) changing the root password

E) configuring secure shell user equivalency for the cellmonitor user

https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmso/exadata-introduction.html

NO.28 You are planning the monitoring configuration for your Exadata X9M Database Machine.

Which three components are monitored directly through the use of Exadata-specific Enterprise Manager Plug-Ins?
* ASM instances
* Hybrid Columnar Compression (HCC) ratios on Extended (XT) storage servers
* the storage server ILOM
* the Power Distribution Units (PDUs)
* Oracle clusterware on the database server
* the RDMA over Converged Ethernet (RoCE) switches
According to Oracle’s documentation12, the components that are monitored directly through the use of Exadata-specific Enterprise Manager Plug-Ins are:

The storage server ILOM2

The Power Distribution Units (PDUs)2

The RDMA over Converged Ethernet (RoCE) switches3

NO.29 Which three steps are required to expand an Exadata X9M-2 Database Server with the lowest memory configuration available to the highest memory configuration?
* Add 12x 32GB DIMMs.
* Add 12x 64GB DIMMs.
* Shutdown the Database Server if running.
* Add 16x 32GB DIMMs.
* Add 16x 64GB DIMMs.
* Memory cannot be expanded on Exadata X9M-2 Database Servers.
* Remove existing memory modules.
* Add 32x 64GB DIMMs.
* Add 24x 32GB DIMMs.

NO.30 You are planning the monitoring configuration for your Exadata X9M Database Machine.

Which three components are monitored directly through the use of Exadata-specific Enterprise Manager Plug-Ins?
* ASM instances
* Hybrid Columnar Compression (HCC) ratios on Extended (XT) storage servers
* the storage server ILOM
* the Power Distribution Units (PDUs)
* Oracle clusterware on the database server
* the RDMA over Converged Ethernet (RoCE) switches

1. The Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book (page 6-25) recommends using the Exadata-specific Enterprise Manager Plug-Ins to monitor ASM instances.

2. Additionally, the Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book (page 6-25) recommends using the Exadata-specific Enterprise Manager Plug-Ins to monitor the storage server ILOM.

3. The Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book (page 6-25) also recommends using the Exadata-specific Enterprise Manager Plug-Ins to monitor Oracle clusterware on the database server.

NO.31 Which three statements are true about the CELLCLI command?
* It can execute commands on multiple storage servers in parallel.
* It has command-line history.
* It requires root privileges to create CELLDISKS and GRIDDISKS.
* It can be executed on storage servers.
* It can be executed using the DCLI utility.

The CellCLI utility supports command-line history.

Each Exadata Storage cell can be monitored, configured, and maintained using the cellcli command line interface.

The dcli utility facilitates centralized management across an Oracle Exadata Storage Server Software realm by automating the execution of CellCLI commands on a set of cells and returning the output to the centralized management location where the dcli utility was run.

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

NO.32 Which are two valid reasons for executing an X9M-2 Exadata storage server rescue procedure?
* the failure of physical disk 1
* corruption in the / (root) filesystem
* the failure of both physical M.2 disks
* the failure of physical disk 0 and 11
* moving all disks from one cell to another as part of a chassis-level component failure
* accidental loss of all data from all griddisks in a storage server
* corruption in a normal or high redundancy ASM diskgroup

The rescue procedure is necessary when system disks fail, the operating system has a corrupt file system, or there was damage to the boot area. If only one system disk fails, then use CellCLI commands to recover. In the rare event that both system disks fail simultaneously, you must use the Exadata Storage Server rescue functionality provided on the Oracle Exadata Storage Server Software CELLBOOT USB flash drive.

NO.33 Which two statements are false about backup to ZFS Storage Appliance (ZFSSA)?
* ZFSSA may connect directly to the Exadata 100Gb RoCE network switches.
* Top of Rack (ToR) switches are managed as part of the hardware stack and software levels are upgraded during the patching process.
* When backing up multi-rack systems, sharing Exadata X9M switches is recommended.
* ZFS Snapshots can provide rapid cloning of development and test environment.

1. ZFS Storage Appliance (ZFSSA) can’t connect directly to the Exadata 100Gb RoCE network switches. ZFSSA uses Fibre Channel protocol to connect to the Exadata storage cells. C. When backing up multi-rack systems, sharing Exadata X9M switches is not recommended. Each rack should have its own switches to minimize the risk of data loss due to switch failure.

NO.34 Which two of the following network connection types can be deployed on the client network ports of an Exadata X9M-2 Database Server?
* 2x 25Gb ports on eth1 and eth2 using SFP28 transceivers and OM4 fiber cable for backup network, and 2x 25Gb ports on eth5 and eth6 using SFP28 transceivers and OM4 fiber cable for client network
* 2x 25Gb ports on eth1 and eth2 using SFP28 transceivers and OM4 fibre cable for client network, and 2x 10Gb ports on eth9 and eth10 using RJ45 cat6 cable for backup network
* 2x 10Gb ports on eth1 and eth2 using FU45 cat6 cable for client network, and 2x 25Gb ports on eth11 and eth12 using SFP28 transceivers and OM4 fiber cable for backup network
* 2x 25Gb ports on eth1 and eth2 using SFP28 transceivers and OM4 fiber cable for backup network, and 2x 10Gb ports on eth3 and eth4 using RJ45 cat6 cable for client network

https://www.oracle.com/a/ocom/docs/engineered-systems/exadata/exadata-x9m-2-ds.pdf

NO.35 An Exadata storage server physical disk on an X9M-2 high-capacity full rack entered the predictive failure state. Which two steps must you perform to replace this failed physical disk?
* Replace the failed physical disk.
* Add the griddisks back into the ASM diskgroup they used to be a member of.
* Identify the griddisks located on the failed physical disk and drop them from the associated ASM diskgroups.
* Rebalance the data on the effected griddisks before performing a manual drop command.
* Create a new celldisk and new griddisks on the replaced physical disk.
* Verify that the griddisks located on the physical disk have been successfully dropped from the associated ASM diskgroups.

Once the physical disk has been replaced, you will need to create a new celldisk and griddisks on the replaced physical disk. After the griddisks have been dropped from the associated ASM diskgroups, you can then add the griddisks back into the ASM diskgroup they used to be a member of and rebalance the data on the effected griddisks. Finally, you should verify that the griddisks located on the physical disk have been successfully dropped from the associated ASM diskgroups.

NO.36 You are adding a disk expansion kit to a running Exadata X8M Database Machine’s Database Servers, and have a filesystem layout that includes:

```
Filesystem                          Mounted on
/dev/mapper/VGExaDb-LVDbSys1        /
/dev/mapper/VGExaDb-LVDbVar1        /var
/dev/mapper/VGExaDb-LVDbHome        /home
/dev/mapper/VGExaDb-LVDbTmp         /tmp
/dev/mapper/VGExaDb-LVDbVarLog      /var/log
/dev/mapper/VGExaDb-LVDbOra1        /u01
/dev/mapper/VGExaDb-LVDbVarLogAudit /var/log/audit
```

After running the following commands, which command needs to be run to add 20G of space to the filesystem mounted on /u01?

```
# parted -s /dev/sda mkpart primary 240132160s 8189439966s
# parted -s /dev/sda set 3 lvm on
# lvm pvcreate --force /dev/sda3
# lvm vgextend VGExaDb /dev/sda3
```

* # lvextend -L +20G -verbose /dev/mapper/VGExaDb-LVDbOral
* # xfs_growfs /uOl +20G
* # resize2fs +20G /dev/VGExaDb/LVDbOral
* # lvextend -L +20G -verbose /dev/VGExaDb/LVDbOral

After running the commands above, the filesystem mounted on /u01 is on the logical volume /dev/mapper/VGExaDb-LVDbOral. So, to add 20G of space to the filesystem mounted on /u01, the command that needs to be run is:

lvextend -L +20G –verbose /dev/mapper/VGExaDb-LVDbOral

This command will extend the logical volume /dev/mapper/VGExaDb-LVDbOral by 20 GB of space. It is important to note that the option –verbose is used to display the progress of the operation.

NO.37 You are in the process of upgrading your nonvirtualized X9M-2 Database Machine elastic configuration with 4 database servers and 7 HC storage servers with an additional 4 database servers and 7 HC storage servers.

The new storage servers are called DM01CEL08 through dmoicel14.

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 DATE diskgroup by executing this SQL statement?

* 1 consisting of all 96 griddisks
* 96 consisting of one griddisk each

* 0 because the new griddisks will be added to the existing faiigroups
* 12 consisting of seven griddisks each
* 7 consisting of 12 griddisks each

This SQL statement is adding the new griddisks to the existing diskgroup “DATA” and creating one new failgroup, consisting of all 96 griddisks. The “REBALANCE POWER 512” option tells the system to perform a rebalance operation with a power of 512. It means the system distributes the data evenly across all the disks in the diskgroup using a power of 512.

NO.38 Which two statements are correct about adding an additional database server to a physical Exadata X9M Database Machine using Oracle Exadata Deployment Assistant (OEDA)?

* Do not proceed if the OEDA Validate Configuration File step displays an error message about missing files p6880880.zip.
* Executing /opt/oracle.supportTools/reclaimdisks.sh -free -reclaim on each Exadata X9M Database server is required to reclaim disk space and perform partition reconfiguration.
* In order to configure the servers with Oracle Exadata Deployment Assistant (OEDA), the new server information must be entered in OEDA, and the configuration file must contain existing nodes.
* The applyElasticConfig.sh script performs network configuration for the new servers. The new servers are restarted at the end of the process.
* It is required to install OEDA on the first new database server.

It is not necessary to execute /opt/oracle.supportTools/reclaimdisks.sh -free -reclaim on each Exadata X9M Database server, nor is it required to install OEDA on the first new database server. In order to configure the servers with Oracle Exadata Deployment Assistant (OEDA), the new server information must be entered in OEDA, and the configuration file must contain existing nodes.

Search results: [1] Refer to the Oracle Exadata Database Machine X9M Implementation Essentials documentation for more information on adding an additional database server to a physical Exadata X9M Database Machine using Oracle Exadata Deployment Assistant (OEDA). [2] Oracle Exadata Deployment Assistant (OEDA) is a graphical user interface (GUI) tool used to configure the Oracle Exadata Database Machine X9M. [3] This article provides the step-by-step instructions to add an additional database server to a physical Exadata X9M Database Machine using Oracle Exadata Deployment Assistant (OEDA). [4] It is important to run the OEDA Validate Configuration File step prior to adding the new server to ensure the configuration file is valid and that all required files exist. [5] When adding a new database server with OEDA [1], the configuration file must contain the existing nodes and the new server information must be entered in OEDA. [6] The applyElasticConfig.sh script must be executed to configure the new servers, after which the new servers will be restarted.

NO.39 Which four actions should you take before proceeding with applying updates to your Exadata Database Machine?

* Consult My Oracle Support note 888828.1 to determine the current recommended Exadata software release.
* Check the Exadata Critical Issues My Oracle Support note 1270094.1 for any issues not added to the latest version of exachk.
* Run exachk and resolve only WARNINGS that you have not seen before.
* Run the appropriate patchmgr preqequisite check step for each component being updated.
* Run patchmgr –all_comp -autofix -autobackup -upgrade -rolling.
* For database servers, perform a server backup using patchmgr -dbnodes db_list_file -backup -rolling.

NO.40 Which three statements are true about the CELLCLI command?

* It can execute commands on multiple storage servers in parallel.
* It has command-line history.
* It requires root privileges to create CELLDISKS and GRIDDISKS.
* It can be executed on storage servers.
* It can be executed using the DCLI utility.

The statements that are true about the CellCLI command are:

It has command-line history . According to 1, you can use the up and down arrow keys to scroll through previous commands that you entered in CellCLI . You can also use Ctrl+R to search for a previous command by entering part of it .

It can be executed on storage servers . According to 1, you can invoke CellCLI by logging in to the Exadata cell as cellmonitor , celladmin , or root user and typing “cellcli” . Within CellCLI , you can perform various management and monitoring tasks for the cell134.

It can be executed using the DCLI utility . According to 5, DCLI (distributed CLI) is another utility that can be used to execute scripts and commands across multiple storage cells from a single interface. You can use DCLI to run CellCLI commands on multiple cells at once54.

NO.41 Which two statements are true about applying updates on Exadata systems?
* Failed storage server updates are automatically rolled back to a previous release.
* To speed up applying storage server updates in a rolling manner, updating two storage servers simultaneously is recommended.
* Updating kernel and RDMA packages on storage cells should be prevented by excluding them with the yum -exclude option.
* For regular Exadata updates, yum automatically installs a non-UEK kernel, which can be selected to boot from grub.
* Failed database server updates are rolled back to a working state on a previous release automatically.
* When running a “yum update” for a new Exadata release, all other repositories should be disabled.
1. Exadata updates have a rollback feature in case the update fails, this means that the system will roll back to a previous release automatically, this applies to storage servers updates. D. Yum, the package manager used in Exadata, automatically installs a non-UEK kernel when running a regular update, this kernel can be selected to boot from grub.

NO.42 I/O performance of the prod database on your Exadata Database Machine has degraded slightly over the past month. The database has been allocated to the OLTP I/O Resource Management (IORM) category.

Which two monitoring tools might be useful in examining I/O performance for the prod database?
* OS I/O metrics using Enterprise Manager host pages for the storage servers
* OS I/O metrics using OS tools such as iostat on the database servers
* I/O-specific dynamic performance views such as v$iostat_fiie, v$iostat_function, and v$iostat_consumer_group from the prod database instances using SQL *p1us
* cellcli (or exacli/exadcli) to examine storage server metrics such as database, category, ceiidisk, and griddisk
* OS I/O metrics using OS tools such as iostat on the storage servers
1. OS I/O metrics using Enterprise Manager host pages for the storage servers: Enterprise Manager provides a detailed view of the storage server’s performance including I/O metrics and statistics. You can use these metrics to analyze I/O performance and identify any bottlenecks.

2. cellcli (or exacli/exadcli) to examine storage server metrics such as database, category, ceiidisk, and griddisk: These commands are used to gather information about storage server performance, including I/O statistics, disk usage, and resource allocation.

NO.43 Which three of the following options are available for database consolidation on Exadata?
* Multiple pluggable databases in one or more container databases on a Virtual Machine or Bare Metal cluster
* Bare Metal deployment with one database per database server
* A single database with one or more schema per application
* Multiple databases spanning a Bare Metal and Virtual Machine cluster to provide resource isolation and consolidation concurrently
* Multiple Databases on a cluster
* Multiple Virtual Machine clusters supporting a single database for greatest resource isolation
According to Oracle.com documents or resources, the three options that are available for database consolidation on Exadata are:

1. Multiple pluggable databases in one or more container databases on a Virtual Machine or Bare Metal cluster123.

2. A single database with one or more schema per application12.

3. Multiple Databases on a cluster12.

https://blogs.oracle.com/exadata/post/database-consolidation-why-and-how

NO.44 Which are three customer options for hosting the Platinum Services Advanced Support Gateway?
* Install on Oracle Database Appliance.
* Provide individual x86 64-Bit gateway hardware.
* Install in Oracle Virtual Machine with required hardware.
* Purchase the recommended x86 64-Bit gateway hardware from Oracle.
* Install on Exadata Engineered System.

According to Oracle’s documentation1, the customer options for hosting the Platinum Services Advanced Support Gateway are:

Provide individual x86 64-Bit gateway hardware1

Purchase the recommended x86 64-Bit gateway hardware from Oracle1

Install on Oracle Database Appliance2

NO.45 You want to monitor how a large production table is accessed. Especially, you are interested to see how the access on that particular table leverages the benefits of the Exadata Platform.

Which two actions are NOT appropriate for that purpose?
* YOU query VSSYSTEM_EVENTS and filter for the event ‘cell physical IO interconnect bytes returned by smart scan’, associated to your table.
* You query v$segment_statistics and filter for the Object ID of your table from dba_objects and the the column STATISTIC_NAME=’optimized physical reads’.
* You query v$SYSTAT and filter for the statistic ‘cell smart table scan’, associated to your table.
* You run the CellCli-command list activerequest , filtering for the attributes ioReason and objectNumber, that you specify as ‘Smart Scan’ and the Object ID of your table from DBA_OBJECTS.

1. Querying VSSYSTEM_EVENTS for the event ‘cell physical IO interconnect bytes returned by smart scan’ will not be appropriate for monitoring the access on a particular table. The V$SYSTEM_EVENT view is used to monitor system-wide performance statistics, not specific table access. C. Querying v$SYSTAT for the statistic ‘cell smart table scan’ will not be appropriate for monitoring the access on a particular table as well. The v$SYSTAT view contains system-wide performance statistics, not specific table access.

NO.46 Which two statements are true about applying updates on Exadata systems?
* Failed storage server updates are automatically rolled back to a previous release.
* To speed up applying storage server updates in a rolling manner, updating two storage servers simultaneously is recommended.
* Updating kernel and RDMA packages on storage cells should be prevented by excluding them with the yum -exclude option.
* For regular Exadata updates, yum automatically installs a non-UEK kernel, which can be selected to boot from grub.
* Failed database server updates are rolled back to a working state on a previous release automatically.
* When running a “yum update” for a new Exadata release, all other repositories should be disabled.

When applying updates on Exadata systems, database server updates are automatically rolled back to a working state on a previous release if the update fails. This ensures that the system can be restored to a working state in the event of a failed update. Additionally, storage server updates are also automatically rolled back to a previous release in the event of a failed update.

This is documented in the Oracle Exadata Database Machine X9M Implementation Essentials official text book and more information can be found on the Oracle support website [1][2].

NO.47 Which statement is true about the patching features provided in Platinum Services?
* Oracle Platinum Service covers Exadata storage software and firmware patching, but customers must perform the database patching.
* Patching services are available for the full software stack up to twice per year.
* The rolling and complete down time approaches are two options to patching.
* Patching is done automatically during business hours.