Easy To Download Oracle 1z1-902 Exam Dumps Updated 60 Questions [Q17-Q34

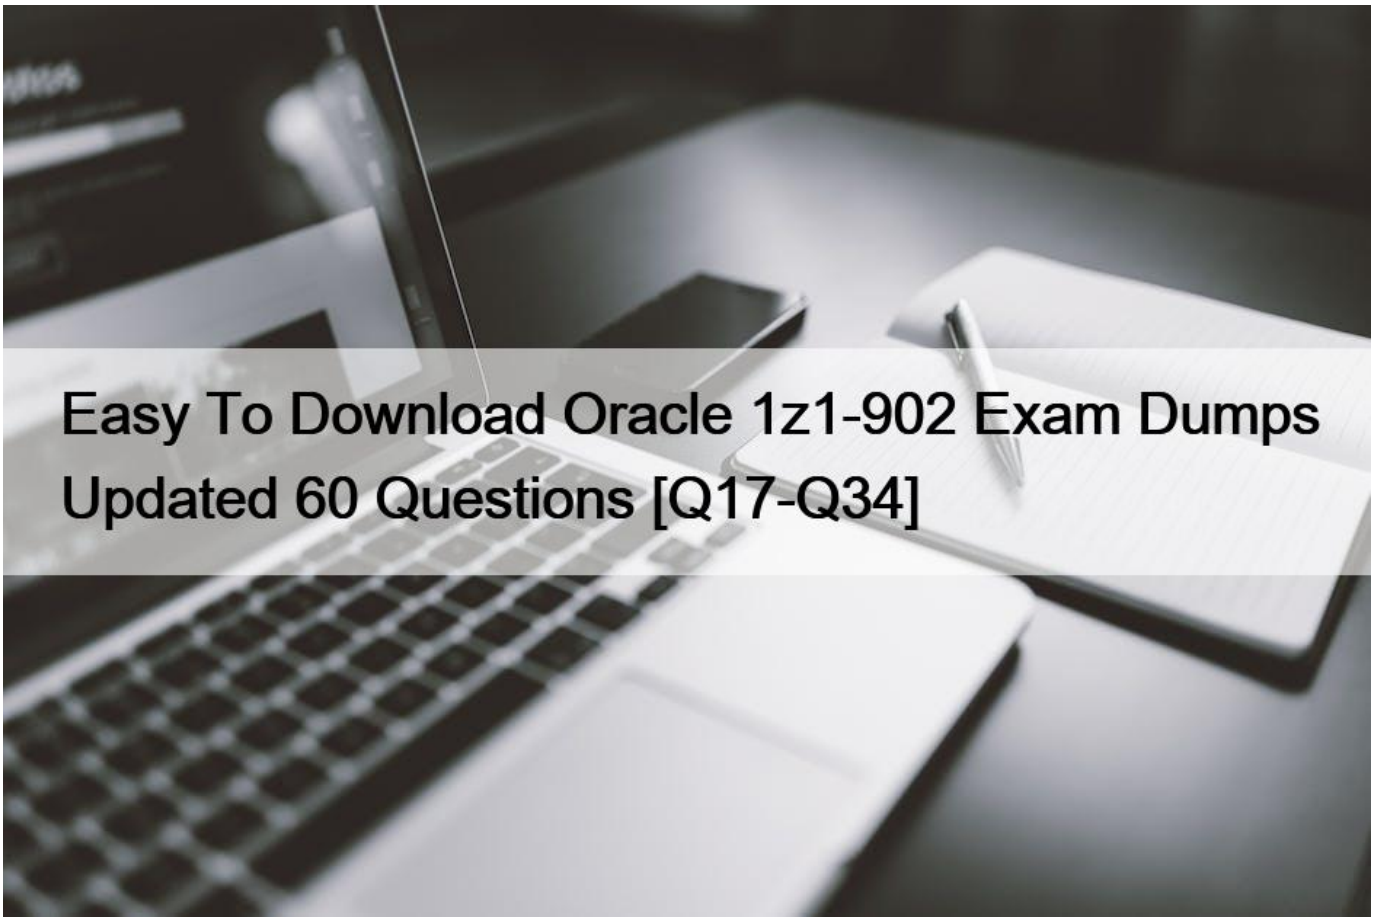

Easy To Download Oracle 1z1-902 Exam Dumps Updated 60 Questions
New Updated 1z1-902 Exam Questions 2023

Oracle 1z1-902 (Oracle Exadata Database Machine X8M Implementation Essentials) Exam is designed to test the knowledge and skills of individuals who are interested in implementing the Oracle Exadata Database Machine X8M. 1z1-902 exam is targeted towards professionals who have experience in database administration, storage administration, and networking. 1z1-902 exam focuses on various topics such as Exadata storage, networking, backup and recovery, and Exadata software.

NO.17 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
According to the Oracle documentation1, two monitoring tools that might be useful in examining I/O performance for the prod database are:

cellcli (or exacli/exadcli) to examine storage server metrics such as database, category, ceiidisk, and griddisk (D). This tool can help you monitor the I/O Resource Management (IORM) metrics and identify any bottlenecks or imbalances in the storage layer.

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 *plus. These views can help you monitor the I/O activity and latency at the file, function, and consumer group level.

NO.18 Which of the following is NOT a requirement when validating, receiving, unpacking, and planning access route and space requirements for Exadata Database Machine?
* The entire access route to the installation site should be free of raised-pattern flooring that can cause vibration.
* 914mm of space required above the rack height is required for maintenance access.
* The incline of any access route ramp must be less than or equal to 6 degrees.
* All four leveling and stabilizing feet should be raised and out of the way prior to moving the rack.
* Oracle Exadata Rack may only be installed on raised floor environments.
* A conditioned space is required to remove the packaging material to reduce particles before entering the data center.
Exadata Database Machine is a pre-configured and pre-tuned hardware and software system designed to run Oracle Database, it can be installed on raised floor environments, but also on concrete or tile floors Oracle Exadata Database Machine X9M Implementation Essentials states that Exadata racks are designed to be installed on flat surfaces, not raised floor environments. It is not required to install the rack on raised floor environments. Additionally, the other options listed are all requirements for validating, receiving, unpacking, and planning access route and space requirements for Exadata Database Machine. (Source: Oracle Exadata Database Machine X9M Implementation Essentials, page 41)

NO.19 Which two options can be used to identify a damaged or failing flash card on an X9M-2 Database Machine High Capacity storage server?
* using the CELLCLI CALIBRATE command on the storage server after logging in as the celladmin user
* using the CELLCLI CALIBRATE command on the storage server after logging in as the root user
* hardware monitoring using the storage server ILOM
* using the CELLCLI list LUN DETAIL command as the celladmin user
The CELLCLI CALIBRATE command can be used to check the health of all flash cards in the storage server, and it should be run after logging in as the celladmin user. The hardware monitoring using the storage server ILOM should also be used to check for errors and other potential issues with the flash cards. The CELLCLI list LUN DETAIL command should not be used to identify a damaged or failing flash card, since it will not provide any information about the health of the flash card.

https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmmn/index.html

NO.20 Your customer needs to ensure that their data is available on the Exadata machine during updates. The customer wants to be able to update one server at a time but still be protected against single-node server failure.

What ASM redundancy level should they use?
* Normal
* Sparse
* High
* External
* Extended

https://www.oracle.com/technetwork/database/exadata/maa-exadata-asm-cloud-3656632.pdf

NO.21 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.

NO.22 You are hardening the security posture of your Exadata Database Machine. Before disabling ssh access to the storage servers, what should you do to enable REST access to the MS process?
* Install Oracle Rest Data Services on each Database server and install the MS APEX application.
* The MS Process on the storage servers is natively endowed with REST services, but are not enabled by default.
* Install Oracle Rest Data Services on each Storage server and install the MS APEX application.
* The MS Process on the storage servers is natively endowed with REST services and are enabled by default, however, appropriate roles and users should be created to ensure security.

The MS process on Exadata storage servers is natively endowed with REST services and these services are enabled by default. However, it is important to create appropriate roles and users to ensure security when accessing these services. This way, it is possible to disable ssh access to the storage servers, and still enable REST access to the MS process. It’s not necessary to install Oracle Rest Data Services on each Database or Storage server or install the MS APEX application, the MS process on the storage servers already have REST services by default.

NO.23 You have been asked to design a backup solution for an Exadata X9M-2 Quarter Rack with Extreme Flash Storage Servers connected to a new ZFS Storage Appliance ZS7 with 2 Storage Controllers with 100Gb Ethernet cards and 3 Storage Trays. You are using Oracle Exadata Configuration Assistant to validate the rack layout.

1. Use “Add Equipment” to add the Exadata X9M EF Storage Servers, starting from RU10.

2. Use drop down to add ZFS Storage Appliance Controllers.

3. You cannot add ZFS Storage Appliance to an Exadata Rack.

4. Use “Add Equipment” to add the ZFS Storage Trays, starting from RU1.

5. Use drop down to add ZFS Storage Trays.

6. Use “Add Equipment” to add the Exadata X9M Database Servers, starting from RU16.

7. Use “Add Equipment” to add the Exadata X9M EF Storage Servers, starting from RU1.

8. Use drop down to add Exadata X9M EF Storage Servers.

9. Use “Add Equipment” to add the ZFS Storage Controllers, starting from RU27.

10. Use drop down to add Exadata X9M Database Servers.

11. Use “Add Equipment” to add the ZFS Storage Trays, starting from RU31.

Which of these steps are correct and what is their correct order?
* 10, 8, 2, 5
* 4,1, 6, 9
* 3
* 10, 8, 9, 11
* 7,6,9,11

The correct order of steps is 10, 8, 9, 11. The 10th step is to use the drop down to add Exadata X9M Database Servers, the 8th step is to use the drop down to add Exadata X9M EF Storage Servers, the 9th step is to use “Add Equipment” to add the ZFS Storage Controllers, and the 11th step is to use “Add Equipment” to add the ZFS Storage Trays. These steps are referenced in the Oracle Exadata Database Machine X9M Implementation Essentials Official Textbook, which is available online at https://docs.oracle.com/cd/E80437_01/E80437/html/index.html.

https://docs.oracle.com/en/engineered-systems/exadata-database-machine/dbmin/configuring-exadata.html

NO.24 Your customer needs to ensure that their data is available on the Exadata machine during updates. The customer wants to be able to update one server at a time but still be protected against single-node server failure.

What ASM redundancy level should they use?
* Normal
* Sparse
* High
* External
* Extended

https://www.oracle.com/technetwork/database/exadata/maa-exadata-asm-cloud-3656632.pdf According to the Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book, the customer should use the High redundancy level for their ASM storage in order to ensure that their data is available on the Exadata machine during updates. The High redundancy level provides the most protection against single-node server failure, as it stores three copies of each data file on separate nodes. This is the same redundancy level used for the Oracle Database’s Control Files and Redo Logs. Instructions for configuring the ASM redundancy level can be found in the book and can be referenced here: https://docs.oracle.com/cd/E80920_01/E80920/html/x9m_asm_overview.html.

NO.25 Which two quarantine types can disable Smart Scan for multiple databases that offload SQL statements to a cell on an Exadata Database Machine?
* SQL Plan Quarantine
* Manually created Quarantine
* Database Quarantine
* Disk Region Quarantine
* Cell Offload Quarantine

A and E are the two correct quarantine types that can disable Smart Scan for multiple databases that offload SQL statements to a cell on an Exadata Database Machine. A is correct because SQL Plan Quarantine will disable Smart Scan for all queries related to the SQL plan that was placed in the SQL Plan Quarantine [1]. E is correct because the Cell Offload Quarantine will disable Smart Scan for all queries offloaded to Oracle Database Exadata Storage Server Software [2]. The other statements (B, C, and D) are incorrect.

[1] Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book , Chapter 13 [1][2]: Oracle Database Exadata Storage Server Software [2] Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book , Chapter 15 [1][2]: Oracle Database Exadata Storage Server Configuration

https://docs.oracle.com/en/engineered-systems/exadata-database-machine/sagug/exadata-storage-server-software-introduction.html

NO.26 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.

According to Oracle.com documents or resources, the four actions that you should take before proceeding with applying updates to your Exadata Database Machine are:

1. Consult My Oracle Support note 888828.1 to determine the current recommended Exadata software release123.

2. Check the Exadata Critical Issues My Oracle Support note 1270094.1 for any issues not added to the latest version of exachk453.

3. Run exachk and resolve only WARNINGS that you have not seen before3.

4. Run the appropriate patchmgr preqequisite check step for each component being updated3.

NO.27 Which two options can be used to identify a damaged or failing flash card on an X9M-2 Database Machine High Capacity storage server?
* using the CELLCLI CALIBRATE command on the storage server after logging in as the celladmin user
* using the CELLCLI CALIBRATE command on the storage server after logging in as the root user
* hardware monitoring using the storage server ILOM
* using the CELLCLI list LUN DETAIL command as the celladmin user

The CELLCLI CALIBRATE command can be used to check the health of all flash cards in the storage server, and it should be run after logging in as the celladmin user. The hardware monitoring using the storage server ILOM should also be used to check for errors and other potential issues with the flash cards. The CELLCLI list LUN DETAIL command should not be used to identify a damaged or failing flash card, since it will not provide any information about the health of the flash card.

NO.28 Which of the following is NOT a requirement when validating, receiving, unpacking, and planning access route and space requirements for Exadata Database Machine?
* The entire access route to the installation site should be free of raised-pattern flooring that can cause vibration.
* 914mm of space required above the rack height is required for maintenance access.
* The incline of any access route ramp must be less than or equal to 6 degrees.
* All four leveling and stabilizing feet should be raised and out of the way prior to moving the rack.
* Oracle Exadata Rack may only be installed on raised floor environments.
* A conditioned space is required to remove the packaging material to reduce particles before entering the data center.

Exadata Database Machine is a pre-configured and pre-tuned hardware and software system designed to run Oracle Database, it can be installed on raised floor environments, but also on concrete or tile floors

NO.29 You have been asked by Oracle Support to check the firmware of the hard & flash disks in a storage server. Which cellcli command should you use to get this information?
* list celldisk attributes name, deviceName, diskType, physicalFirmware
* list cell attributes name, deviceName, diskType, physicalFirmware
* list physicaldisk attributes name, deviceName, diskType, physicalFirmware
* list devices attributes name, deviceName, diskType, physicalFirmware

NO.30 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.31 You have been asked to investigate why an Exadata Database Server stopped communicating on the client network for 10 minutes over the past weekend.

Which command would help investigate this?
* $ ${ORACLE_HOME}/suptools/tfa/release/tfa_home/bin/tfact1 -from <Fridays_Date>_17:00:00 -to <Sundays_Date>_23:59:00
* # /opt/oracle.SupportTools/ibdiagtools/netcheck/runDiagnostics.pm -from <Fridays_Date>_17:00:00 -to <Sundays_Date>_23:59:00
* # /opt/oracle.ExaWatcher/GetExaWatcherResults.sh -from <Fridays_Date>_17:00:00 -to <Sundays_Date>_23:59:00
* # /opt/oracle.SupportTools/exachk/exachk -from <Fridays_Date>_17:00:00 -to <Sundays_Date>_23:59:00

Exachk is an Oracle-provided diagnostic tool that allows you to check various components of an Exadata Database Server, including the client network. The tool is located in the /opt/oracle.SupportTools/exachk/ directory and can be run with the command: exachk. This command can be run with the option -from <Fridays_Date>_17:00:00 and -to <Sundays_Date>_23:59:00 to specify a time range of interest, this way you can check the client network status between the specified time range. The tool checks the status of various components such as network interfaces, InfiniBand fabric, and storage servers. It also analyzes log files, configuration files, and system settings to detect any issues. The tool generates a report, which can be used to identify the root cause of the issue and provide recommendations for resolution.

Tfact1 is a tool to collect diagnostic information from Exadata and non-Exadata systems and it is not related to client network status.

Netcheck and GetExaWatcherResults.sh are not tools provided with the Exadata Database Server and they are not related to client network status.

NO.32 Which dbmcli command is NOT valid on Exadata X9M?
* dbmcli -e “LIST METRICHISTORY WHERE name LIKE ‘DS_.*'”

B) dbmcli -e “LIST METRICCURRENT WHERE name = ‘DS_TEMP’ “
* dbmcli -e “LIST IBPORT DETAIL”
* dbmcli -e “LIST ALERTHISTORY WHERE agelnMinutes < 15”

NO.33 Which two statements are true about the initial storage configuration after the standard (non-virtualized) deployment of a new Exadata Database Machine with High Capacity storage servers?
* The sparse_<DBM_NAME> diskgroup is created automatically.
* There is free space available on the hard disks inside the database servers for possible extension of the /uoi file system.
* The DATA_<DBM_Name> and RECO_<DBM_NAME> ASM diskgroups are built on with DATA on the outer-most tracks and RECO on the inner-most tracks of the physical disk.
* There is free space available on flashdisks inside the Exadata storage servers for possible use for storage indexes.
* There is free space available on flashdisks inside the Exadata storage servers to configure Exadata Smart Flash Logs.

According to the Oracle Exadata Database Machine Technical Architecture1, the initial storage configuration after the standard (non-virtualized) deployment of a new Exadata Database Machine with High Capacity storage servers includes two ASM disk groups: DATA_<DBM_Name> and RECO_<DBM_NAME>. These disk groups are built on hard disks and flash disks inside the Exadata Storage Servers1.

The correct statements about this configuration are:

Option D: There is free space available on flashdisks inside the Exadata storage servers for possible use for storage indexes. Storage indexes are a feature of Exadata Storage Software that can improve query performance by avoiding unnecessary I/O operations. Storage indexes use a small amount of flash memory to store metadata about data blocks stored on disk1.

Option E: There is free space available on flashdisks inside the Exadata storage servers to configure Exadata Smart Flash Logs. Exadata Smart Flash Logs are another feature of Exadata Storage Software that can improve database performance by using flash memory as an extension of the database redo log buffer. This can reduce latency and increase throughput for redo log writes1.

NO.34 You are providing oversight for the delivery of a new Exadata Database Machine.

1. Stabilize the Exadata Rack.

2. Unpack Oracle Exadata Rack.

3. Review the safety guidelines.

4. Let the Exadata acclimatize for 24 hours.

5. Power on Exadata PDU A.

6. Place Exadata in its allocated space.

7. Power on Exadata PDU B.

What is the correct order of these steps?
* 4,3,2,6,1,7,5
* 3,2,6,4,1,7,5
* 2,3,4,6,1,7,5
* 3,2,6,1,4,7,5
* 2,6,1,4,3,7,5

This ordering is based on the Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book and other resources. First, the Exadata should be stabilized and unpacked (Steps 2 and 3). Then, the safety guidelines should be reviewed (Step 4). The Exadata should then be allowed to acclimatize for 24 hours (Step 4). After that, the Exadata PDU A should be powered on (Step 6) and the Exadata should be placed in its allocated space (Step 6). Finally, the Exadata PDU B should be powered on (Step 7).