[Full-Version 2025 New Preparation Guide of Oracle 1Z0-902 Exam [Q42-Q61

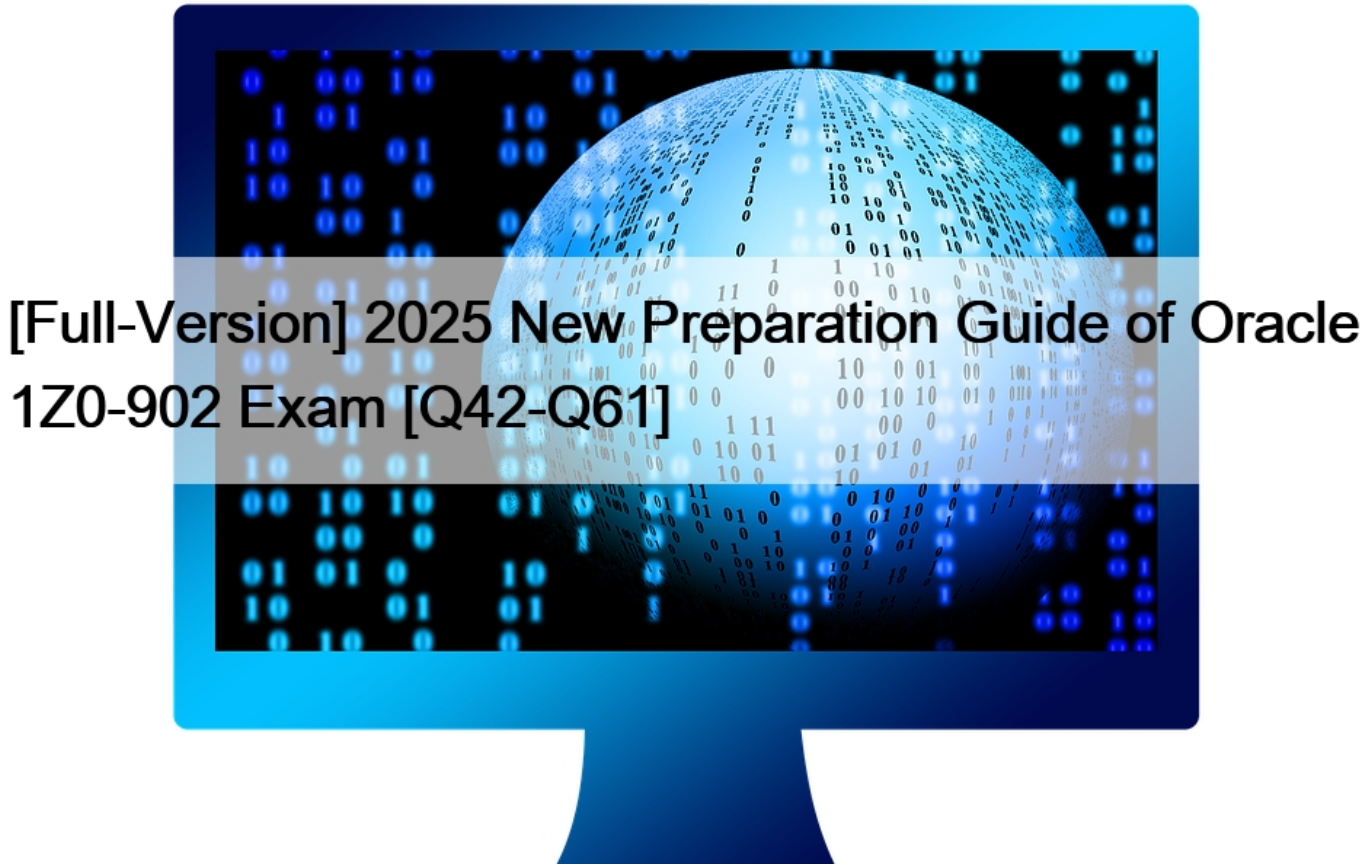

[Full-Version] 2025 New Preparation Guide of Oracle 1Z0-902 Exam
1Z0-902 Practice Exam - 87 Unique Questions

QUESTION 42

How is Exadata Secure RDMA Fabric isolation used in X9M-2 with virtualization enabled?
* With Secure Fabric isolation, each Storage Server runs multiple KVM guests each dedicated to a Database VM cluster using a dedicated network partition and VLAN ID. This provides security and isolation between multiple tenants or VM clusters.
* With Secure Fabric isolation, each Database VM Guest uses a dedicated network partition and VLAN ID for client networking between application servers. This improves security by preventing man in the middle attacks.
* With Secure Fabric isolation, each Database VM Cluster uses a dedicated network partition and VLAN ID for Data Guard networking between the primary and disaster recovery Exadata racks. This improves security by isolating data guard traffic onto an encrypted network between data centers.
* With Secure Fabric isolation, each Database VM cluster uses a dedicated network partition and VLAN ID for cluster networking between the database servers in the VM Cluster. This provides security and isolation between multiple tenants or VM clusters.
Exadata Secure RDMA Fabric isolation is a feature that allows you to use virtualization on the X9M-2 platform while maintaining

security and isolation between multiple tenants or VM clusters. Secure Fabric isolation works by creating a dedicated network partition and VLAN ID for each Database VM cluster, which is used for cluster networking between the database servers in the VM Cluster. This ensures that the network traffic of each VM cluster is isolated from the others, and it prevents unauthorized access to the data or resources of other VM clusters. Secure Fabric isolation can also be used for other purposes such as client networking, Data Guard traffic, but it’s mainly used to provide security and isolation between multiple tenants or VM clusters.

According to Oracle’s documentation1, Exadata Secure RDMA Fabric Isolation is a feature that enables strict network isolation for Oracle Real Application Clusters (Oracle RAC) clusters on Exadata Database Machine systems that use RDMA over Converged Ethernet (RoCE). Exadata Secure RDMA Fabric Isolation uses RoCE VLANs to ensure that network packets from one VM cluster cannot be seen by another VM cluster. VLAN tag enforcement is done at the KVM host level, which means that security cannot be bypassed by any software exploits or misconfiguration on the database server VMs1.

Therefore, the statement that is true about how Exadata Secure RDMA Fabric isolation is used in X9M-2 with virtualization enabled is:

With Secure Fabric isolation, each Database VM cluster uses a dedicated network partition and VLAN ID for cluster networking between the database servers in the VM Cluster. This provides security and isolation between multiple tenants or VM clusters.

QUESTION 43

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

QUESTION 44

Which two statements are true in regards to starting the Exadata Virtual Machine?
* To see Oracle Linux boot messages during guest startup, use -console option with the vm_maker -start-domain command.
* Use vm_maker -start-domain to start a virtual machine manually.
* Use vm_maker boot-from-iso command to boot a virtual machine.
* To streamline the diagnosis of virtual machines, one ISO file is used for multiple Oracle Exadata System Software releases.
* Use vm_maker -auto-start command to configure a virtual machine to start automatically when the KVM host is started.

QUESTION 45

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.

QUESTION 46

Which two statements are true about the IPTables firewall configuration on a new Exadata Database Machine after initial deployment?
* IPTables is installed and available but not configured or running on the storage servers.
* IPTables is configured and running with Oracle-supplied rules on the storage servers.
* IPTables is configured and running with Oracle-supplied rules on the database servers.
* IPTables is installed and available but not configured or running on the database servers.
* IPTables is installed and available but not configured on any server.

On Oracle Exadata Database Machine X9M, the default configuration of the IPTables firewall varies between storage servers and database servers:

Storage Servers: After initial deployment, IPTables is installed but remains in an unconfigured and inactive state. The Exadata storage servers focus on providing high availability and performance for database workloads and thus rely on internal security measures. The external firewall or network security policies typically handle firewall responsibilities, so IPTables is not actively

configured on these servers by default.

Database Servers: Similar to storage servers, IPTables is also installed but neither configured nor running on the database servers. The Oracle Exadata Database Machine X9M emphasizes security through other means, such as its network isolation and dedicated interconnects. As part of Oracle’s engineered systems design philosophy, the database servers are primarily focused on database performance, offloading much of the security management to other parts of the infrastructure and relying on data center-level firewalls for perimeter security.

To summarize, both database and storage servers have IPTables installed and available, but they are not configured or running by default on either server type. This setup allows Exadata environments to focus on optimized database performance while integrating with external network security solutions.

QUESTION 47

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

A Exadata supports multiple pluggable databases in one or more container databases, which enables database consolidation by allowing multiple databases to be consolidated into a single container database. This can be deployed on a Virtual Machine or Bare Metal cluster.

D Multiple databases spanning a Bare Metal and Virtual Machine cluster provide resource isolation and consolidation concurrently. It allows multiple databases to be deployed on both Bare Metal and Virtual Machine clusters, providing isolation and consolidation at the same time.

E Multiple Databases on a cluster allows multiple databases to be deployed on a single cluster, which enables consolidation of databases by reducing the number of clusters required.

QUESTION 48

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.

https://docs.oracle.com/en/database/oracle/oracle-database/21/ladbi/db-on-exadata.html

QUESTION 49

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

QUESTION 50

A new Exadata Quarter Rack with 2 Database Servers and 3 HC Storage Servers and 3-phase 15kVA PDUs is being installed in a Data Center. However, the Data Center is only providing enough power for a single cable from each PDU.

Which statement is correct?
* The installation can go ahead, no change is required.
* The installation cannot proceed until two power feeds are available per PDU.
* A splitter cable can be used to provide power to all PDU cables.
* The power cables from the servers to the PDUs can be rearranged inside the rack following OECA guidance to utilize a single PDU power cable.

An Exadata Quarter Rack with 2 Database Servers and 3 HC Storage Servers and 3-phase 15kVA PDUs would typically require two power feeds per PDU, one for each of the two power distribution units (PDUs) in the rack. A single power feed per PDU is not sufficient to provide the necessary power to the servers and storage in the rack.

QUESTION 51

Examine this list of software components:

1. Oracle KVM Guest

2. Oracle Enterprise Manager Agent (OMA)

3. ASM instance

4. RDBMS instance

5. Automatic Diagnostic Repository Command Interpreter (ADRCI)

6. CELLCLI

7. Cell Server(CELLSRV)

8. diskmon

9. Restart Server (RS)

10. Management Server (MS)

What is the correct location where these software components can run in the standard Exadata Database Machine deployment?
* 2, 3, 4, 8, and 10 run on the database servers; 1, 5, 6, 7and 9 run on the Exadata storage servers.
* 1, 2, 3, 4, 9 and 10 run on the database servers; 5, 6, 7, 8, 9, and 10 run on the Exadata storage servers.
* 1, 2, 3, 4, 5, 8, 9 and 10 run on the database servers; 5, 6,7, 9 and 10 run on the Exadata storage servers.

* 3, 4, 8, and 10 run on the database servers; 1, 2, 5, 6, 7 and 9 run on the Exadata storage servers.
* 1, 2, 3, 4, 8 and 9 run on the database servers; 5, 6, 7, 9 and 10 run on the Exadata storage servers.

Oracle KVM Guest, Oracle Enterprise Manager Agent (OMA), RDBMS instance, Automatic Diagnostic Repository Command Interpreter (ADRCI), CELLCLI, diskmon, and Restart Server (RS) can all run on the database servers. The Cell Server (CELLSRV) and Management Server (MS) can both run on the Exadata storage servers. Specific instructions for installing and configuring these software components can be found in the book and can be referenced here: https://docs.oracle.com/cd/E80920_01/E80920/html/x9m_software_components.html.

QUESTION 52

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 two valid reasons are:

corruption in the / (root) filesystem . According to 1, this is one of the scenarios that requires a storage server rescue procedure , because it prevents the storage server from booting up normally . The rescue procedure will re-image the root filesystem with a fresh copy of the operating system and restore the configuration files from a backup1.

the failure of both physical M.2 disks . According to 2, each Exadata Storage Server X9M-2 has two physical M.2 disks that store the operating system and configuration files for the storage server . If both M.2 disks fail , then the storage server cannot boot up or function properly . The rescue procedure will replace the failed M.2 disks with new ones and re-image them with the operating system and configuration files2.

QUESTION 53

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

QUESTION 54

Which are two correct statements about backing up Exadata Database Machine?
* Backup of the Oracle Cluster Registry, which also contains Voting Disk information, is automatically maintained on the file system of the first database server.
* Backing up Exadata Storage Server Software is critical for restoration.
* For high availability, system area can be on the first two disks and M.2 devices.

* The M.2 disk is hot pluggable and can be replaced when the power is on, but you will have to restore the system data manually.
* Use a snapshot-based backup of an Oracle Exadata Database Machine database server software.

QUESTION 55

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.
This is according to the Oracle Exadata Database Machine X9M Implementation Essentials Official Text Book [1]. The other statements are false as there is no free space available on the hard disks inside the database servers for possible extension of the /uoi file system, and there is no free space available on flashdisks inside the Exadata storage servers for possible use for storage indexes or to configure Exadata Smart Flash Logs [2][3].

QUESTION 56

You are working on a remote Exadata Database Machine with current software and firmware. You now need to start a storage server after hardware maintenance.

Which ILOM command starts the storage server?
* -> start /System
* -> start /SP/console
* -> set /SYS power_state=ON
* You must be physically present in the data center to start an Exadata Storage Server.
* -> start /SYSTEM/power
* -> start /SYS/power

QUESTION 57

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.

QUESTION 58

Which two statements are true about Auto Service Request (ASR) with an Exadata Database Machine?
* ASR Manager must be installed and configured on a dedicated server external to the Database Machine.
* Configuring ASR is mandatory for all Database Machine assets.
* ASR Manager must be installed and configured on one of the database servers.
* ASR can upload configuration metadata to support problem resolution.
* ASR Manager opens a service request (SR) automatically after sensors detect hardware faults.
* ASR communicates with Oracle support services using HTTPS.

QUESTION 59

Examine these commands:

1. Execute “crscti stop cluster -all” as the grid user from one database server.

2. Execute “crscti stop cluster -all” as root from one database server.

3. Power off all network switches.

4. Execute “crscti stop cluster” as root from one database server.

5. Execute “crscti stop cluster” as the grid user from one database server.

6. Power off the rack using the power switches on the PDUs.

7. Execute “shutdown -h now” on all database servers.

8. Execute “shutdown -h now” on all Exadata storage servers.

Which is the correct order or the required commands to completely power off an Exadata Database Machine in an orderly fashion?
* 5, 8, 7, and 6
* 4, 7, 8, 3, and 6
* 2, 8, 7, 3, and 6
* 2, 7, 8, and 6
* 1, 8, 7, 3 and 6

In order to power off an Exadata Database Machine in an orderly fashion, the commands should be executed in the following order: 1) Execute “crscti stop cluster -all” as the grid user from one database server; 2) Execute “shutdown -h now” on all database servers; 3) Power off all network switches; 4) Execute “shutdown -h now” on all Exadata storage servers; and 5) Power off the rack using the power switches on the PDUs. This sequence of commands ensures that all services running on the Exadata Database Machine are in a clean state before powering off the system. (Source: Oracle Exadata Database Machine X9M Implementation Essentials, page 543)

QUESTION 60

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)

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

QUESTION 61

Your system administrator reports an amber, non-blinking light on one of your Exadata storage disks. You immediately execute the ‘list physicaldisk where diskType=HardDisk and status=failed DETAIL’ command on the Exadata storage system and the specified disk is indeed reported as failed. Platinum Support has not been enabled for this system. Enterprise Manager Cloud Control is monitoring the system.

What is the next step that you should perform before you do anything else?
* Ask the system administrator to replace the broken disk with a spare.
* Check the database to see if any rebalance operations are active.
* Call Oracle Support and make an appointment so that the drive can be replaced.
* Wait for the email of the failure that Exadata or Enterprise Manager will send.
* Download and run the latest exadiag tool.
* Wait for a blue light to appear on the disk if the rebalance operation is running.