1Z0-058^{Q&As}

Oracle Real Application Clusters 11g Release 2 and Grid Infrastructure Administration

Pass Oracle 1Z0-058 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.passapply.com/1Z0-058.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle Official Exam Center

🔧 **Instant Download** After Purchase

🔧 **100% Money Back** Guarantee

🔧 **365 Days** Free Update

🔧 **800,000+** Satisfied Customers

QUESTION 1

The database administrator is tasked with creating an ASM disk group. Exadata is not being used.

If failure groups are not specified when creating an ASM disk group containing 10 disks, how many failure groups are automatically created?

A. one

B. two

C. five

D. ten

Correct Answer: D

There are always failure groups even if they are not explicitly created. If you do not specify a failure group for a disk, then Oracle automatically creates a new failure group containing just that disk, except for disk groups containing disks on Oracle Exadata cells. Oracle?Automatic Storage Management Administrator\\'s Guide

QUESTION 2

Which statement describes the requirements for the network interface names, such as eth0, in Oracle Clusterware?

A. Only the public interface names must be the same for all nodes.

B. Only the private interface names must be the same for all nodes.

C. Both the public interface name and the private interface name must be the same for all nodes.

D. Both the public interface name and the private interface name can vary on different nodes.

E. Only the private interface names can be different on different nodes.

F. Only the public interface names can be different on different nodes.

Correct Answer: C

Checking Network Requirements

Each node must have at least two network interface cards (NICs).

Interface names must be the same on all nodes.

Public NIC must support TCP/IP and Private NIC UDP.

Public IP must be registered in the domain name server (DNS) or the /etc/hosts file.

cat /etc/hosts

Public Interfaces eth0 (odd numbers)####

xxx.xxx.100.11 host01.example.com host01

xxx.xxx.100.13 host02.example.com host02 IF GNS is used, the Cluster GNS address must be registered in the DNS. Prevent public network failures when using NAS devices or NFS mounts by starting the Name Service Cache Daemon. # /sbin/service nscd start

D60488GC11 Oracle 11g: RAC and Grid Infrastructure Administration Accelerated 2 18

QUESTION 3

After evaluating the various methods for extending a cluster, you decide to use addNode.sh.

The cluster originally consisted of four nodes: RACNODE1, RACNODE2, RACNODE3, and RACNODE4. Now two nodes called RACNODES and RACNODE6 have been installed and connected to the cluster by OS administrations.

Which three actions should be performed to check whether the new nodes are ready for running addNode.sh and to help correct any problems?

A. cluvfy stage -pre crsinst -n RACNODE5/ RACNODE6 -C + DATA -q +VOTE -orainv

B. -fixup -verbose

C. cluvfy stage -post hwos -n RACNODE5, RACNODE6 -verbose

D. cluvfy comp peer -refnode RACNODE1 -n RACNODE5, RACNODE6 -orainv -osdba -verbose

E. cluvfy stage -post hwos -n all -verbose

F. cluvfy stage -pre nodeadd -n RACNODE5, RACNODE6 -fixup

G. cluvfy comp peer -refnode RACNODES -n RACNODE6 -orainv -osdba -verbose

Correct Answer: CDF

Run the following CVU command to check cluster integrity. This command verifies that any number of specified nodes has been successfully added to the cluster at the network, shared storage, and clusterware levels:

$ cluvfy stage -post nodeadd -n node3 [-verbose]

cluvfy comp peer

Use the cluvfy comp peer component verification command to check the compatibility and properties of the specified nodes against a reference node. You can check compatibil- ity for non-default user group names and for different releases

of the Oracle software. This command compares physical attributes, such as memory and swap space, as well as user and group values, kernel settings, and installed operating system packages.

Syntax

cluvfy comp peer -n node_list [-refnode node] [-r {10gR1 | 10gR2 | 11gR1 | 11gR2}] [- orainv orainventory_group][-osdba osdba_group] [-verbose] Usage Notes Peer comparison with the -refnode option compares the system properties of

other nodes against the reference node. If the value does not match (the value is not equal to reference node value), then CVU flags that comparison as a deviation from the reference node. If a group or user does not exist on reference

node

as well as on the other node, CVU reports this comparison as \\'passed\\' because there is no deviation from the reference node. Simil- arly, CVU reports as \\'failed\\' a comparison with a node that has more total memory than the reference node.

Verify the integrity of the cluster and node3:

$ cluvfy stage -pre nodeadd -n node3 [-fixup [-fixupdir fixup_dir]] [-verbose] You can specify the -fixup option and a directory into which CVU prints instructions to fix the cluster or node if the verification fails.

Oracle?Clusterware Administration and Deployment Guide

QUESTION 4

You are managing a three-Instance RAC database. In the AWR report, you notice a gc current block busy wait event on one of the database Instances, in the cluster top timed events section. What are two possible reasons for this wait event?

A. Access to blocks was delayed as the blocks were pinned In exclusive mode for updates by some sessions.

B. Access to blocks was delayed by log writes on remote Instance.

C. Frequently used select statements are causing high disk Input/output contention.

D. CPU shortages.

Correct Answer: AB

Analyzing Performance Using GCS and GES Statistics

The gc current block busy wait event indicates that the access to cached data blocks was delayed because they were busy either in the remote or the local cache. This could be caused by any of the following:

The blocks were pinned

The blocks were held up by sessions

The blocks were delayed by a log write on a remote instance

A session on the same instance was already accessing a block which was in transition between instances and the current session needed to wait behind it (for example, gc current block busy) Contention-Related Wait Events

The main wait events for contention-related waits are:

gc current block busy

gc cr block busy

gc buffer busy acquire/release

The contention-related wait event statistics indicate that a block was received which was pinned by a session on another node, was deferred because a change had not yet been flushed to disk or because of high concurrency, and therefore

could not be shipped immediately. A buffer may also be busy locally when a session has already initiated a cache fusion operation and is waiting for its completion when another session on the same node is trying to read or modify the same

data. High service times for blocks exchanged in the global cache may exacerbate the contention, which can be caused by frequent concurrent read and write accesses to the same data. The gc current block busy and gc cr block busy wait

events indicate that the local instance that is making the request did not immediately receive a current or consistent read block. The term busy in these events\\' names indicates that the sending of the block was delayed on a remote instance.

For example, a block cannot be shipped immediately if Oracle Database has not yet written the redo for the block\\'s changes to a log file.

Oracle?Real Application Clusters Administration and Deployment Guide 11g Release 2 (11.2)

QUESTION 5

Which two statements are true about ACFS snapshots?

A. They can be created for ACFS file systems only if the ASM disk group hosting the ADVM volume file used by the file system has free space available.

B. They can be created for ACFS file systems only if the ADVM volume file used by the file system has free space available.

C. They can be created only if the ASM disk group hosting the ADVM volume used by the file system has no other ASM files contained in the disk group.

D. They can be created when ACFS is used both on clusters and on stand-alone servers.

E. They are accessible only on the cluster node that was used when creating the snapshot.

Correct Answer: BD

About Oracle ACFS Snapshots Oracle ACFS snapshot storage is maintained within the file system, eliminating the management of separate storage pools for file systems and snapshots. Oracle ACFS file systems can be dynamically resized to accommodate additional file and snapshot storage requirements. Oracle?Automatic Storage Management Administrator\\'s Guide 11g Release 2 (11.2)

[1Z0-058 PDF Dumps](#)   [1Z0-058 Exam Questions](#)   [1Z0-058 Braindumps](#)

To Read the Whole Q&As, please purchase the Complete Version from Our website.

Try our product !

100% Guaranteed Success
100% Money Back Guarantee
365 Days Free Update
Instant Download After Purchase
24x7 Customer Support
Average 99.9% Success Rate
More than 800,000 Satisfied Customers Worldwide
Multi-Platform capabilities - Windows, Mac, Android, iPhone, iPod, iPad, Kindle

We provide exam PDF and VCE of Cisco, Microsoft, IBM, CompTIA, Oracle and other IT Certifications.
You can view Vendor list of All Certification Exams offered:

https://www.passapply.com/allproducts

Need Help

Please provide as much detail as possible so we can best assist you.
To update a previously submitted ticket: