

Oracle 1z0-884

**Oracle Solaris Cluster 3.2 System Administrator
Certified Professional Exam**

Practice Test

Version: 14.20

QUESTION NO: 1

A system administrator has received the Java ES distribution as a two-CD set, and wishes to configure an NFS server to use as an install server for Sun Cluster 3.2.

What is the correct procedure?

- A. Mount and share each CD from a separate server.
- B. This is NOT possible, only the DVD distribution can be used in this case.
- C. Create a working directory, share it using NFS, and combine the contents of both CDs into it using the cpio command.
- D. Mount the first CD and share it using NFS. When prompted by the installer, unshare and unmount the CD, and repeat the process with the second CD.

Answer: C

Explanation:

QUESTION NO: 2

You run `cldg status` to view the status of your device group. What must be true about the node that is listed as the Primary?

- A. It is the preferred node for the device group.
- B. It is the current primary node for the device group, not necessarily the preferred node.
- C. A global file system configured on a volume of the device group can only be accessed from that node.
- D. The device group can run only on that node. If that node fails, the device group will become unavailable.

Answer: B

Explanation:

QUESTION NO: 3

The cluster administrator is running SunCluster 3.2 on x86 machines running Solaris 10 update 3, and needs to boot one of the x86 machines into non-cluster mode.

To which menu item does the cluster administrator add `-x` using the GRUB editor?

- A. kernel /platform/i86pc/multiboot

- B. kernel /platform/i86pc/suserboot
- C. Solaris 10 11/06 s10x_u3wos_08 X86
- D. module /platform/i86pc/boot_archive

Answer: A

Explanation:

QUESTION NO: 4

The SUNW.HAStoragePlus resource manages different types of storage. In which case is it mandatory to configure HAStoragePlus?

- A. to manage global raw devices
- B. to manage global file systems
- C. to manage failover file systems
- D. to manage global file system devices with an Oracle RAC configuration

Answer: C

Explanation:

QUESTION NO: 5

A customer needs to do some maintenance to an application resource in their NFS resource group. They want to have NFS resources active, but want to be sure that if anything goes wrong during the operation, the NFS service is NOT restarted or the resource group switched.

How can this be accomplished ?

- A. The customer can suspend the resource group.
- B. The customer can place the resource group into a quiescent state.
- C. The customer can place the resource group offline and then enable only the required resources.
- D. The customer can keep the resource group online and then disable monitoring for all the resources.

Answer: A

Explanation:

QUESTION NO: 6

Which commands would be used to define a third private network using switches for a two-node cluster?

- A. `clintr add-privnet -s switch3 node1:bge2,node2:bge2`
- B. `clintr add node1:bge2`
- C. `clintr add node1:bge2`
- D. `clintr add -s switch3 node1:bge2 node2:bge2`

Answer: B

Explanation:

QUESTION NO: 7

A two-node Solaris Volume Manager campus cluster has been installed at sites A and B. A preferred site A can be configured so that the cluster survives the loss of site B without requiring manual intervention.

Which is the correct procedure to create a preferred site A ?

- A. Put the quorum device at site A and configure the majority of the diskset metadb replicas on the disks located at site A.
- B. Put the quorum device at site A and add an extra metadb to the disks in the diskset at both sites A and B.
- C. Put the quorum device at site A and add an extra local metadb to the root disk on the server located at site A.
- D. You cannot do this unless you configure a third site with the quorum device and with the diskset configured across all three sites, so that the loss of one site will always ensure that two out of the three metadbs are present.

Answer: A

Explanation:

QUESTION NO: 8

You have two adapters that you want to configure in the same IPMP group. Neither of them are yet configured into any IPMP group. The cluster nodes are already configured.

Which tasks must you perform to configure the group?

- A. You need to rerun the `scinstall` utility and let it create your IPMP groups for you.

- B.** You need to create or edit the `/etc/hostname.xxx` files for each adapter, using the group option to place them in the same group.
- C.** You need to use the `clinterconnect (clintr)` command and specify the adapters for which you want to create IPMP groups.
- D.** You need to create or edit the `/etc/hostname.xxx` file for only one adapter. In that file, you list all the adapters that will be in the same group.

Answer: B

Explanation:

QUESTION NO: 9

A patch has been released for a recently discovered security issue. The cluster administrator has downloaded the patch and the README shows that it is a rebooting node patch that must be installed in single-user mode.

What steps should be followed to apply this patch with the minimum amount of downtime for the service.

A. On each node in turn run the commands:

```
clnode evacuate -n  
shutdown -gO -y -i0  
boot -sx
```

Apply the patch.

Reboot back into cluster.

B. On one node run the following commands:

```
clrg status  
clrg offline +
```

On each node in turn run the commands:

```
shutdown -g0 -y -i0  
boot -sx
```

Apply the patch.

Reboot back into cluster.

When all nodes are patched run the following commands from one node:

```
clrg online +  
clrg status
```

C. Connect to all nodes using the cluster console and apply the patch simultaneously.

Reboot each node in turn.

D. On one node run the commands:

```
clrg status  
cluster shutdown -g0 -y
```

Reboot each node outside cluster.

```
boot -sx
```

Apply the patch on each node and reboot back into cluster.

boot

Answer: A

Explanation:

QUESTION NO: 10

There are various ways to interact with Sun Cluster 3.2 to do daily administration.

Which method is prohibited?

- A. using the clsetup command
- B. editing the ccr
- C. using Sun Cluster Manager
- D. using standard cluster commands

Answer: B

Explanation:

QUESTION NO: 11

You just created a Solaris Volume Manager diskset from the shared storage containing two disks, one from each of two arrays. You added two connected nodes as mediators for the diskset.

Which two of the following are true? (Choose two.)

- A. You must manually create a partition on the diskset disks to hold a metadb replica.
- B. The metainit commands that you use to build volumes on the disks will fail unless you mirror across arrays.
- C. The mediators allow you to lose a disk and subsequently maintain high availability if a node fails.
- D. If you lose one of the disks and then lose the node that is primary for the diskset, the entire cluster will lose access to the diskset.
- E. The Solaris Volume Manager commands that added disks to the diskset automatically created a single metadb replica on each disk.

Answer: C,E

Explanation:

QUESTION NO: 12

Which are three reasons to disable automatic quorum device configuration in a two-node cluster? (Choose three.)

- A. You have a dual ported disk or LUN that is not capable of being a quorum device.
- B. You want to use a second internal drive for your quorum device.
- C. You want to use the quorum server as the quorum device.
- D. You want to use a NAS device as a quorum device.
- E. You want to use your boot device as your quorum device.

Answer: A,C,D

Explanation:

QUESTION NO: 13

What is the minimum number of nodes required to make a highly available cluster?

- A. One
- B. Two
- C. Four
- D. Eight

Answer: B

Explanation:

QUESTION NO: 14

You have a two-node cluster configured. Node one is out of the cluster. You make changes to node two (add resource groups and disk device groups). You now shut down node two and try to boot node one, which fails due to amnesia prevention.

Which two keep node one from starting the cluster? (Choose two.)

- A. scsi-2 reservations
- B. Persistent reservations on the quorum device.
- C. Nothing will keep node one from starting the cluster.
- D. Node one's key has been previously preempted from the quorum device.

Answer: B,D

Explanation:**QUESTION NO: 15**

A customer has a resource group called nfs-rg with a SUNW.nfs resource called nfs-rs. They want to check how the Failover_mode property has been set.

Which is the correct command used to accomplish this task ?

- A. clrg status
- B. clrs show -g nfs-rg
- C. clrs show -v nfs-rs
- D. clrs list -v nfs-rs

Answer: C

Explanation:

QUESTION NO: 16

What is the complete list of acceptable Veritas volumes to hold your Sun Cluster 3.2 application data on arrays of JBOD disks?

- A. mirrored stripe, striped mirrors
- B. simple mirrors, mirrored stripe, striped mirrors
- C. concatenated subdisks, striped subsets, simple mirrors
- D. striped mirrors, striped subdisks, concatenated subdisks, concatenated mirrors

Answer: B

Explanation:

QUESTION NO: 17

You need to have two failover IP addresses on separate subnets both fail over along with an application. How do you create your IP address resources?

- A. You can create a single LogicalHostname resource representing both IP addresses.
- B. You must create two separate LogicalHostname resources and put them in the same resource group as the application resource.
- C. You must create two separate LogicalHostname resources and put them in two separate

resource groups. A resource group can only contain one LogicalHostname resource.

D. You must create a SharedAddress resource. Only a SharedAddress resource can refer to multiple IP addresses.

Answer: B

Explanation:

QUESTION NO: 18

Which two types of communication can occur on the private interconnect? (Choose two.)

- A.** cluster heartbeats
- B.** application data transfer
- C.** vxconfigd communication
- D.** Solaris Volume Manager communication

Answer: A,B

Explanation:

QUESTION NO: 19

You need to remove the Veritas Volume Manager volume named nfsvol in the nfsdg disk group which is mounted on /nfs.

Which set of commands are used to perform this procedure?

- A.** umount /nfs; cldg remove nfsdg -r volume nfsvol; sconfg sync nfsdg
- B.** umount /nfs; vxassist -g nfsdg remove volume nfsvol; cldg sync nfsdg
- C.** unmount -g nfsdg all; vxassist -g nfsdg remove all; cldg sync nfsvol
- D.** unmount /nfs; vxassist -g nfsdg remove volume nfsvol; cldg sync nfsvol

Answer: B

Explanation:

QUESTION NO: 20

You have a three-node cluster and have put one node into maintenance mode.

What do you need to do to give the node its vote back and bring it back into the cluster?

- A. Re-install the node to fix the maintenance mode.
- B. Reboot the node, the maintenance mode will reset automatically.
- C. Remove the node from the cluster, then add it back into the cluster.
- D. Add another quorum disk device before booting the node back into the cluster.

Answer: B

Explanation:

QUESTION NO: 21

Which statement is incorrect regarding the in.mpathd daemon?

- A. IPMP uses ping to contact any router listed in the routing table.
- B. If no router exists, IPMP submits a ping to the all-hosts multicast address.
- C. If static host routes are available, these hosts will be contacted using ping.
- D. You can configure link state monitoring on Solaris 9.

Answer: D

Explanation:

QUESTION NO: 22

Occasionally you may need to rebuild your cluster. If you rebuilt a cluster named "dev" without removing the old quorum server information, you may want to clean up the old quorum server information manually.

Which command would you use to remove the old quorum server information after using the command `clqs show` to obtain the cluster name and ID?

(--- Cluster dev (id 0x448DE82F) Registrations ---)

- A. `clclust remove -F quorum-serv-v`
- B. `clqs clear -c dev -l 0x448DE82F 9000`
- C. `scquorum -c dev -l 0x448DE82F 9000`
- D. Reboot the quorum server

Answer: B

Explanation: