Vendor: Oracle

Exam Code: 1Z0-884

Exam Name: Oracle Solaris Cluster 3.2 System Administrator Certified Professional Exam

Version: Demo

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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 respresenting 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

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

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; scconf 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

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

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

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 -I 0x448DE82F 9000
C. scquorum -c dev -I 0x448DE82F 9000
D. Reboot the quorum server

Answer: B

QUESTION NO: 23

You want to create a role that contains only the authorizations Solaris.cluster.modify and Solaris.cluster.admin. As you then decide which users should be able to successfully access every cluster command, you can just assign these users to the role and give them the role password. These users will never need to have the root password, but they will be able to execute every cluster command.

Which two statements are true? (Choose two.)

A. This scheme will work for both cluster commands and Sun Cluster Manager.
B. There are still some cluster-related activities, such as modifying the /etc/vfstab file, that the user might NOT be able to do.
C. The users who switch to the role will still NOT have Solaris.cluster.read authorization, so they will not be able to view cluster status and configuration.
D. This scheme will work as the users try to access command line commands, but it will NOT work if the users execute operations through Sun Cluster Manager.

Answer: A,B

QUESTION NO: 24

Which two Veritas Volume Manager packages need to be installed on all nodes connected to the shared storage? (Choose two.)

A. VRTSjre
B. VRTSvlic
C. VRTSvxvm
D. VRTSobgui
E. VRTSvxman

Answer: B,C

QUESTION NO: 25

When using ZFS for your shared storage within the cluster, how do you specify the file system mountpoint?

A. ZFS does not use mountpoints.
B. You must define mountpoints in the /etc/vfstab file first.
C. ZFS mountpoints default to /poolname/fsname, but you can change them to whatever you want using the zfs command
D. Sun Cluster chooses the name of the file system mountpoints for ZFS.

Answer: C

QUESTION NO: 26

When using ZFS for your shared storage within the cluster, the mountpoints default to "/poolname/fsname". You can change the mountpoint name to whatever you want. You want to change the mountpoint for the filesystem myfs1 in the pool devpool to /oradb.

Which command do you use?

A. zset mount=/oradb devpool/myfs1
B. zpool setmount=/oradb poolname/fsname
C. zfs /oradb set mountpoint devpool/all
D. zfs set mountpoint=/oradb devpool/myfs1

Answer: D

QUESTION NO: 27

Which is applicable to the Cluster Configuration Repository (CCR)?

A. The CCR is only accessed by the cluster software at boot time.
B. The CCR contains information that is critical to the operation of the cluster.
C. The CCR should be configured on shared storage so that all nodes can access it
D. To ensure consistency between the cluster nodes the CCR should be manually copied to all nodes on a regular basis.

Answer: B

QUESTION NO: 28

Which two components are part of the Sun Cluster 3.2 software stack? (Choose two.)

A. operating system
B. Sun Cluster 3.2 agents
C. Sun Cluster 3.2 framework
D. volume management software

Answer: B,C

QUESTION NO: 29

When using ZFS to maintain your cluster file systems, what other software volume management application do you need to manage the device groups?

A. You do not need any other volume manager software.
B. SVM is needed to manage your disksets.
C. VxVM is needed to manage your disk groups.
D. Veritas CVM is needed to manage the shared disk groups.

Answer: A

QUESTION NO: 30

Solaris 10 allows resource allocation to users by assigning them projects. Consider having to automatically assign a project to an Oracle user, to start your HA-Oracle database.

Which two options would accomplish this task? (Choose two.)

A. Use the newtask command
B. Assign a project to the user using only the /etc/project file.
C. Assign a project to the user using both the /etc/project and the /etc/user_attr file.
D. Use the Resource_project_name property to pass the project to be used when starting Oracle.

Answer: C,D

QUESTION NO: 31

After you install a two-node cluster and reboot, before the cluster administrator assigns a quorum device, what state is the cluster in?

A. bootmode
B. installmode
C. clustermode
D. startupmode

Answer: B

QUESTION NO: 32

When using Veritas Volume Manager in the cluster, why would you choose to use Dirty Region Logging (DRL) for your mirrored volumes?

A. DRLs can only be used for RAID 5 volumes.
B. You would never use DRLs, they are too large.
C. You CANNOT use DRLs within Sun Cluster 3.2.
D. You should use DRLs to minimize the resynchronization of your mirrored volumes after a system crash.

Answer: D

QUESTION NO: 33

Which two statements are true about a global file system? (Choose two.)

A. It is seen as mounted only on one node.
B. All nodes have simultaneous access to the data.
C. All nodes directly connected to the storage can do direct i/o.
D. All nodes do NOT have to be directly connected to the storage.

Answer: B,D

QUESTION NO: 34

You configure Oracle10g RAC in Sun Cluster 3.2 with the data on a global file system. What will happen?

A. You will get an error from the database configuration utility.
B. It will actually configure and run, but it is not a supported Oracle configuration.
C. It will run correctly at first, but will fail if the entire cluster ever reboots.
D. It will seem to configure correctly. The database software will actually check the configuration and return an error at run time.

Answer: B

QUESTION NO: 35

Which two can be performed when using Solaris 10 zones? (Choose two.)

A. Run the complete cluster framework in non-global zones.
B. Configure a highly available nfs server only in a non-global zone
C. Install an application that can fail over between non-global zones.
D. Install an application that can fail over between the global zones and the non-global zones.

Answer: C,D

QUESTION NO: 36

When configuring Veritas Cluster Volume Manager for use with a two-node Oracle RAC Sun Cluster 3.2 install, which node can configure the shared disk group?

A. Both nodes can configure the group.
B. Neither node can configure the group.
C. The master node can configure the group.
D. You cannot use CVM in Sun Cluster 3.2.

Answer: C

QUESTION NO: 37

After the cluster administrator has installed the Sun Cluster 3.2 packages, you configure the Sun Cluster 3.2 framework. This can be done in various ways.

Which configuration method does NOT work?

A. Using clsetup to configure the framework.
B. Use Jumpstart to both, install the packages and configure the cluster.
C. Using scinstall to configure one node at a time.
D. Using scinstall to configure the whole cluster a once.

Answer: A

QUESTION NO: 38

An administrator is deciding whether to use a global file system or a failover file system for an application. Which is a good reason to use a global file system?

A. ZFS (and its new features) can be used as a global file system.
B. A global file system can be accessed by all nodes of the cluster simultaneously.
C. Writes to a global file system can have a performance benefit over failover file systems
D. You can only configure a shared copy of application binaries in a global file system.

Answer: B

QUESTION NO: 39

A new cluster has been installed and the customer needs to create a HA-NFS service. Which steps have to be executed ?

A. Register SUNW.nfs dataservice, create a resource group, and create a SUNW.nfs resource.
B. Register SUNW.nfs dataservice, create a resource group, create a SUNW.HAStoragePlus resource, and create NFS resource.
C. Register SUNW.nfs dataservice, create a resource group, create a SUNW.HAStoragePlus resource, and create NFS resource.
D. Create a resource group, create a SUNW.Logicalhostname resource, create a SUNW.HAStoragePlus resource.

Answer: C

QUESTION NO: 40

Which configuration can be used with ZFS and Sun Cluster 3.2?

A. ZFS as a global file system.
B. ZFS as a failover file system.
C. ZFS as a multi-owner file system.
D. A ZFS file system in /etc/vfstab (setting the file system mount point to legacy' within ZFS).

Answer: B

QUESTION NO: 41

What is the name of the cluster console package that the cluster administrator will load onto the administrative workstation?

A. SUNWsc
B. SUNWccon
C. SUNWsccon
D. SUNWconsole

Answer: B

QUESTION NO: 42

Which tasks are performed by the Sun Cluster 3.2 framework components?

A. The CMM causes resource groups to fail over between cluster nodes.
B. The RGM is used to start and stop the application on the cluster nodes.
C. The CCR manages the cluster-wide agreement on the set of nodes booted into the cluster.
D. If the RGM detects a critical problem with a node, it calls upon the cluster framework to forcibly shut down (panic) the node.

Answer: B

QUESTION NO: 43

How do you decide what DID name to give to your Veritas Volume Manager disks as you assign them to disk groups?

A. The cluster assigns the DID name to the VxVM disk.
B. VxVM automatically assigns the DID name to the disk.
C. You have to manually assign the DID name to a VxVM disk.
D. VxVM uses its own naming scheme, it does NOT use the DID name.

Answer: D

QUESTION NO: 44

Cluster framework daemons are started with boot scripts on Solaris 9 and with SMF on Solaris 10.

Which statement is true?

A. On Solaris 10, you need to manually run svcadm enable for the cluster framework daemons after the cluster is initially configured.
B. On Solaris 10, every time a node boots it will NOTjoin the cluster automatically. You type svcadm enable for aemon to have it join the cluster.
C. On Solaris 9, you can run /etc/init.d/clusterdaemons stop to manually disable the cluster framework on a running node. On Solaris 10 there is no such operation.
D. The daemons start automatically at cluster boot time regardless of whether you are running Solaris 9 or Solaris 10.

Answer: D

QUESTION NO: 45

You have a customer with a two-node cluster. The only storage accessible from both nodes is a NetApp NAS device. What kind of quorum configuration should the customer choose?

A. They must use a NetAPP NAS quorum device only.
B. They must use a quorum server quorum device only.
C. They can choose either the NetAPP NAS quorum device or the quorum server quorum device. If they choose both, they have too many quorum votes.

D. They should configure both a NetAPP NAS quroum device and a quorum server quorum device. This will give them the correct number of total quorum votes and the best redundancy.

Answer: C

QUESTION NO: 46

Before installing Sun Cluster 3.2 software with Solaris Volume Manager on the nodes that will make up the cluster, which three requirements must be met? (Choose three.)

A. All root partitions must be the same size.
B. The /usr file system must be on a separate partition.
C. A /globaldevices partition must be created.
D. A minimum of 750 Mbytes of swap must be configured.
E. Solaris Volume Manager requires a separate partition for its meta-state databases.

Answer: C,D,E

QUESTION NO: 47

What path should a system administrator add to the MANPATH variable to find the manual pages for the commands in the Sun Cluster 3.2 CLI?

A. /usr/share/man
B. /opt/cluster/man
C. /usr/cluster/man
D. /usr/cluster32/man

Answer: C

QUESTION NO: 48

One of the cluster nodes is at the ok boot prom prompt. You want to boot that node to multiuser and into non-cluster mode. Which command will accomplish this?

A. init 6

B. boot -x

C. boot -sx

D. /etc/reboot -- -x

Answer: B

QUESTION NO: 49

If you administer a cluster with HA-Oracle you should understand some principles of the HA-Oracle fault monitor. Which part is NOT performed by the Oracle fault monitor?

A. The HA-Oracle agent does some table manipulation.

B. The HA-Oracle agent checks some Oracle views for database activity.

C. The HA-Oracle agent creates the fault monitor user automatically, if the user is not present.

D. The HA-Oracle agent scans the Oracle Alert log for specific Oracle errors.

Answer: C

QUESTION NO: 50

A customer has a three-node cluster configured as a pair-plus-one topology. The customer asks about whether they need to use any quorum devices, and if so how many.

Which answer best describes how the administrator will respond?

A. You do NOT need any quorum devices.

B. You must have a quorum device, all clusters require quorum devices.

C. You do NOT need any quorum devices, although it would make your cluster more highly available. You should configure ONE quorum device for highest availability

D. You do NOT need any quorum devices, although it would make your cluster more highly available. You should configure TWO quorum devices for highest availability.

Answer: D

QUESTION NO: 51

Every resource in Sun Cluster 3.2 has to have a resource type. Some of the resource types are pre-registered, and most of them you have to register.

Which resource type listed below is pre-registered?

A. SUNW.gds
B. SUNW.nfs
C. SUNW.LogicalHostname
D. SUNW.HAStoragePlus

Answer: C

QUESTION NO: 52

A customer called saying their cluster is NOT starting. The customer explained that they halted node two for maintenance, then made some modification to the cluster configuration on node one. Then node one crashed. They tried to boot node two but the node is NOT allowed to form a cluster.

What is the result?

A. Split brain occurred. Check private network connection.
B. Cluster amnesia was being prevented. The customer needs to boot node one before node two.
C. Failure fencing is preventing node two from booting in order to preserve data integrity.
D. Disconnect the private network in order to disable node one, which is preventing node two from booting.

Answer: B

QUESTION NO: 53

In Veritas Volume Manager 4.x or 5.x, once you encapsulate your boot device it goes into a disk group. This group historically was called rootdg.

What is the name of the link to this diskgroup under 4.x and 5.x control?

A. rootdg
B. bootdg

C. There is no link to this diskgroup.
D. You must link this group to rootdg.

Answer: B

QUESTION NO: 54

You set up your cluster using the default 172.16.x.x range for the cluster transport. Yet these same IP addresses are in use on machines in a department that you need to access across the public network.

What kind of problem will you experience?

A. Your cluster nodes will refuse to communicate across the private network.
B. Your public network adapters will be marked as FAILED.
C. Communication to the machines using 172.16.x.x on the public network fails, because your cluster nodes wil want to route this traffic across the private network.
D. You will remain free of problems. It is perfectly acceptable to use a range of IP addresses on your private networks that conflicts with other machines on your public network.

Answer: C

QUESTION NO: 55

Which two statements identify the correct boot device restrictions? (Choose two.)

A. There must be a minimum 512 Mbytes /globaldevices file system.
B. There must be a minimum 512 Gbytes /globaldevices file system.
C. You can use a boot device that is located in a SAN, if it is only accessible from that node.
D. You can use a shared storage device as a boot device. If a storage device is connected to more than one host, it is shared.

Answer: A,C

QUESTION NO: 56

You have a ZFS pool containing a single disk c2t0d0. The pool contains a file system that you want to mount on /data.

What does the entry look like in the /etc/vfstab file?

A. ZFS populates the /etc/vfstab entries for you automatically
B. using ZFS, you do NOT create /etc/vfstab entries
C. /dev/dsk/zpool1/dev/rdsk/zpool1 /data zfs 3 no -
D. /dev/dsk/c2t0d0s0 /dev/rdsk/c2t0d0s0 /data zfs 1 yes global

Answer: B

QUESTION NO: 57

Which is a reason for using link-state testing rather than "ping testing" for public network adapters in IPMP groups in Solaris 10?

A. Link-state testing is always more robust.
B. IPMP will do repair detection and fallback only if you use link-state testing.
C. The cluster will only be able to monitor the state of IPMP groups if you use link state testing
D. You do NOT have to consume public network IP addresses for test addresses if you use link-state testing.

Answer: D

QUESTION NO: 58

Solaris 10 uses SMF as a standard method of managing services.

Which two are applicable? (Choose two.)

A. Each SMF service has its own log file for diagnosing startup problems.
B. An application can be made to fail over by configuring an SMF probe within Sun Cluster.
C. SMF gives the administrator the ability to restart Sun Cluster without having to reboot the server.
D. SMF-specific proxy agents can be used to turn SMF-based applications into clustered applications.

Answer: A,D

QUESTION NO: 59

When using ZFS in the cluster environment, you generally do not need any other volume manager to manage your data storage. Instead, you create pools of storage to manage the file systems within.

What is the ZFS storage unit called?

A. zraid
B. zstor
C. zpool
D. zunit

Answer: C

QUESTION NO: 60

Which two are valid reasons to disable automatic quorum configuration? (Choose two.)

A. To configure a two-node cluster.
B. To use a NAS device as a quorum device.
C. To automatically reset the installmode flag.
D. To use the quorum server as a quorum device.
E. To have several dual-ported disks or LUNs configured.

Answer: B,D

QUESTION NO: 61

If the cluster administrator wants to restart a process, they have to be sure it is NOT one of the cluster daemons.

Which two daemons listed below can you send a kill -HUP signal to without affecting Sun Cluster

3.2? (Choose two.)

A. rgmd
B. rpc.pmfd
C. inetd
D. syslogd
E. rpc.fed

Answer: C,D

QUESTION NO: 62

When configuring the RAC framework Resource Group for an Oracle RAC database instance, what do the parameters of Maximum_primaries=, Desired_primaries= refer to?

A. A RAC Framework Resource Group is only for HA Oracle.
B. These are NOT valid options for a RAC Framework Resource Group.
C. The desired number of nodes allowed to run, and the maximum number of nodes allowed in the cluster.
D. The desired number of instances you would like to run, and the maximum number of instances allowed to run.

Answer: D

QUESTION NO: 63

How can you name Solaris Volume Manager disksets that include disks from shared storage?

A. You can name them anything you want, often related to what application you may be using them for.
B. The name of the diskset must be the same as the name of the cluster.
C. The Solaris Volume Manager software automatically chooses a name for the diskset, such as ds101 or ds102. You must keep this name.
D. A diskset that is physically attached to nodes named vincent and theo must be named ds_vincent_theo, and there can be only one such diskset.

Answer: A

QUESTION NO: 64

In Sun Cluster 3.2, what command is used to view the status of IPMP group members on all nodes?

A. scstat -IPMP
B. IPMP status -v
C. clnode status -m
D. clrg public -v all

Answer: C

QUESTION NO: 65

Which two statements identify the tasks a node can perform, when that node has imported a Veritas disk group? (Choose two.)

A. The node will increase the size of the private region on the disks within the disk group.
B. The node physically reads and writes data to the drives within the disk group.
C. The node waits until all mirrored data is in sync before writting data to the disk group.
D. The node can voluntarily give up ownership of the disk group by deporting the disk group.

Answer: B,D

QUESTION NO: 66

Which two options in /etc/vfstab are valid for the global file system? (Choose two.)

A. Mount at boot is set to yes.
B. Mount options must be set to false.
C. The mount point must have the word global.
D. The word global is one of the mount options.

Answer: A,D

QUESTION NO: 67

While setting up logical network interfaces within IPMP groups, what is the maximum number of virtual network interfaces for each physical adapter?

A. You may have only two virtual interfaces for each physical adapter.
B. You may have an unlimited number of virtual interfaces for each physical adapter.
C. You can have a maximum of 8192 virtual interfaces for each physical adapter.
D. You can have as many virtual interfaces as you have physical nodes in the cluster.

Answer: C

QUESTION NO: 68

Before you can configure your cluster, what must you do on each node to prepare for the global devices device files?

A. You do not do anything.
B. You must have unused space on the SAN.
C. Global devices come from your root partition.
D. You must have an unused partition on your boot device.

Answer: D

QUESTION NO: 69

What are the requirements if you want to use only IPV6 test address for your public network adapters, and no IPV4 test address?

A. You cannot use IPV6 for the IPMP interfaces.
B. You can use IPV6, but you must have at least one IPV4 test address.
C. You must not have IPV6 configured on any servers outside the cluster.
D. You must have at least one server outside the cluster configured using IPV6.

Answer: D

QUESTION NO: 70

With Sun Cluster 3.2 it is possible to specify typical or custom installation.

Which two items must you provide to scinstall when performing a typical configuration? (Choose two.)

A. cluster name
B. /globaldevices size
C. cluster transport adapters
D. quorum disk name
E. cluster interconnect network address and netmask
F. number of cluster interconnects to be used and switch configuration

Answer: A,C

QUESTION NO: 71

You are required to integrate an Oracle database into the cluster framework using the HA-Oracle agent. You have configured a logical host and failover file system for the database on shared storage with a mount point of/u01 and need to integrate the database into the cluster framework.

You enter the following command on node1:

/usr/cluster/bin/clrs create -g ora-rg -t SUNW.oracle_server -p ORACLE_SID=SID1

-p ORACLE_HOME=/u01 /oracle/product/10.2.0/db_1

-p Alert_log_file=/u01/oracle/admin/SID1/bdump/alert_SID1.log

-p parameter_file=/u01/oracle/admin/SID1/pfile/initSID1.ora -p connect_string=sc_fm/sc_fm ora-server-res and get the following output:

Mar 13 21:04:25 node2 SC[SUNWscor.oracle_server.validate]:ora-rg:ora-server-res:Validation failed. ORACLE binaries not found ORACLE_HOME=/u01/oracle/product/10.2.0/db_1

What is required to successfully register the agent?

A. The message indicates that the dba has NOT yet created the Oracle database.
B. The Oracle binaries cannot be installed on a failover file system and should be reinstalled on local storage on each cluster node.
C. The ora-server-res resource should be made dependent on the failover file system with the -p Resource_dependencies property.

D. The Oracle listener resource type should be registered first so that the cluster framework can connect to the database with the connect_string sc_fm/sc_fm to verify the database is running.

Answer: C

QUESTION NO: 72

When you design the network redundancy for a cluster node, you have to understand the key requirements of IPMP. Which statement is incorrect?

A. Every adapter must have a unique local MAC address.
B. A network adapter can be a member of only one group.
C. Solaris 9 requires test addresses for every IPMP group with more than one adapter.
D. Solaris 10 must have test addresses configured in an IPMP group with more than one adapter.

Answer: D

QUESTION NO: 73

Sun Cluster 3.2 can take advantage of the new ZFS file system available with Solaris 10.

Which two new features are available with ZFS and Sun Cluster 3.2? (Choose two.)

A. ZFS has a built-in volume manager.
B. ZFS can be used as a server replication method.
C. The disks are used as a storage pool.
D. ZFS can be used as a highly available global file system.

Answer: A,C

QUESTION NO: 74

A customer has a two-node cluster with one quorum server. They have shut down node1 for maintenance. Later in the day they bring down node2 and the quorum server.

The customer has now decided to bring the cluster back up. Which answer best describes the

proper technique?

A. Clear the quorum server keys, then boot node2.
B. Clear the quorum server keys, then boot node1.
C. Boot the quorum server first and then node2.
D. Boot the quorum server first and then node1.

Answer: C

QUESTION NO: 75

Which two name service lookup entries are modified with an entry of cluster in /etc/nsswitch.conf
for Sun Cluster 3.2? (Choose two.)

A. hosts
B. group
C. passwd
D. ethers
E. netmasks

Answer: A,E

QUESTION NO: 76

During a custom scinstall when using switches for your private transport you need to assign them
names.

What names must you use?

A. It does NOT matter what the names are.
B. The switches do NOT require names.
C. Sun gives you the switch name.
D. The switch names must match the interface names.

Answer: A

QUESTION NO: 77

A system administrator has a two-node cluster to which it becomes necessary to add a third node. Once the hardware has been connected, which two steps are required to add the new node to the cluster? (Choose two.)

A. Shut down and reconfigure all cluster nodes.
B. Run scinstall on the third node to add it to the existing cluster.
C. Use the cluster set command to modify the cluster properties.
D. On one existing cluster node use the clsetup or claccess command to add the new hostname to the authorized node list.

Answer: B,D

QUESTION NO: 78

Which two are the default settings for Disk Path Monitoring? (Choose two.)

A. all disk paths are monitored
B. only disk paths to shared devices are monitored
C. reboot_on_path_failure is enabled
D. reboot_on_path_failure is disabled

Answer: A,D

QUESTION NO: 79

What is required to do before installing the Solaris (SUNWudlm) and Oracle (ORCLudlm) packages?

A. Create an Oracle database.
B. Create the Oracle dba group.
C. Configure the RAC framework in the cluster.
D. Configure the shared storage on which the RAC database will be created.

Answer: B

QUESTION NO: 80

Which IPMP restriction for multiple adapters in an IPMP group is no longer a restriction in Solaris 10?

A. Every adapter must have a dedicated test IP address.
B. A network adapter can be a member of only one group.
C. Each Ethernet adapter must have a unique MAC address.
D. All members of the same group must be on the same subnet.

Answer: A

QUESTION NO: 81

You want to create separate non-global zones on different nodes, and be able to fail over an application between the two zones. What would be the advantage of using the same zone name on different nodes?

A. You can set up failover between non-global zones on different nodes only if those zones have the same name
B. You can use the syntax -n node1,node2 -z zonename only if the zone has the same name on both nodes.
C. The zones can get simultaneous access to data that is in a global file system only if they have the same name.
D. If the zones have the same name, then you will be able to switch the resource group to the zone on a specific node by simply typing clrg switch -n nodename rgname

Answer: B

QUESTION NO: 82

A three-node cluster contains three resource groups which have these Nodelists:

one-rg has a Nodelist: node1 node2 node3

two-rg has a Nodelist: node2 node3 node1

three-rg has a Nodelist: node3 node1 node2

The Fallback property on each of the resource groups is set to FALSE.

Following a maintenance outage in which node1 and node2 have been shut down and rebooted into the cluster with node1 joining the cluster before node2:

Where will the resource groups be located, and which command should be run to move the resource groups to their preferred nodes?

A. one-rg will be on node1
two-rg will be on node2
three-rg will be on node3
No command will need to be run.
B. one-rg will be on node3
two-rg will be on node3
three-rg will be on node3
clrg remaster +
C. one-rg will be on node1
two-rg will be on node1
three-rg will be on node3
clrg remaster +
D. one-rg will be on node3
two-rg will be on node3
three-rg will be on node3
clrg evacuate -n node3 +

Answer: B

QUESTION NO: 83

When designing the hardware for a multi-node cluster, you should care about the hardware redundancies of each node. Some of the possible hardware redundancies are absolutely required, and some are recommended.

Which hardware redundancy is recommended but NOT required?

A. mirrored boot disks
B. redundant transports
C. redundant server nodes
D. multiple paths from each connected node assuming a hardware RAID for shared storage

Answer: A

QUESTION NO: 84

The Sun Cluster 3.2 framework software has already been installed. In order to support a cluster-unaware application, which additional software must be installed?

A. data service agent
B. cluster framework
C. Resource Group Manager
D. Solaris Operating System

Answer: A

QUESTION NO: 85

When installing Sun Cluster 3.2 using the Java ES installer, two configuration options are offered:

Configure Now

Configure Later

In what circumstances is it appropriate to select the first option, Configure Now?

A. If Solaris 10 is already fully installed and patched.
B. If the system being installed is a single-node cluster.
C. When installing a completely new cluster for the first time.
D. Never, Sun Cluster 3.2 must always be configured separately after installation of the packages.

Answer: D

QUESTION NO: 86

During the scinstall process there are many steps executed automatically.

Which two execute automatically when using scinstall on a Solaris 9 system? (Choose two.)

A. creation of the /etc/notrouter file
B. configuration of the resource groups
C. configuration of the volume management
D. selection of a quorum device in a three-node cluster
E. modification of the nsswitch.conf file with the necessary values

Answer: A,E

QUESTION NO: 87

You want to configure five disksets within a cluster environment running on Solaris 10. What must you do for Solaris Volume Manager to allow you to create more than four disksets?

A. Nothing is required, the number of disksets is dynamically updated in Solaris 10.
B. The maximum number of disksets that can be configured in Solaris 10 is four.
C. Each diskset must be registered into the cluster environment using cldg sync
D. Adjust the parameter md_nsets in /kernel/drv/md.conf to allow more than four disksets to be created.

Answer: A

QUESTION NO: 88

When you choose a transport adapter while running scinstall, you are asked:

Why would you answer no?

A. Because you want to use the same physical adapter to hold the traffic for two VLAN's, both on the private net.
B. Because your adapter will be used only for the cluster transport.
C. Because your adapter will use the tagged VLAN feature to support both private and public network traffic.
D. Because your adapter will not be dedicated to the cluster heartbeat, but will also support global file system traffic between nodes.

Answer: C

QUESTION NO: 89

Suppose you have a Veritas Volume Manager disk group named appdg with default settings registered using:

#cldg create -t vxvm -n node1,node2 appdg.

You are now asked to configure the following: appdg should be online on node1 by default, and node1 should take appdg automatically back when it joins the cluster.

Which identifies the correct action to accomplish this task?

A. No action is required, since this is already the default setup.
B. cldg set -p preferenced=true -p failback=true appdg
C. cldg set -p preferenced=node1 -p failback=true appdg
D. cldg change -p preferenced=true -p failback=true appdg

Answer: B

QUESTION NO: 90

Suppose you have two nodes called node1 and node2, and each of them has a non-global zone running called zone1.

Which two commands register a failover resource group named apache-rg to be potentially mastered on both non-global zones? (Choose two.)

A. clrg create -z zone1 apache-rg
B. clrg create -n node1,node2 -z zonel apache-rg
C. clrg create -n zone1@node1,zone1@node2 apache-rg
D. clrg create -n node1:zone1,node2:zone1 apache-rg

Answer: B,D

QUESTION NO: 91

You have a four-node cluster, where node1 and node2 are connected to some shared storage,

and node3 and node4 are connected to different shared storage.

Now node1 has died and you will not be able to reboot it for some time. Why would you then want to run clq disable node1?

A. This will prevent node1 from accidentally rebooting until you explicitly run clq enable node1 on one of the remaining nodes.
B. You may want to run the command but it is actually too late. You can only run the disable option when the node mentioned is still up in the cluster.
C. This will alter the quorum vote count in the remaining cluster so that if later you also lost node2, the remaining nodes could continue running the cluster.
D. This is simply an informational flag. The output of the clq status will specifically show a maintenance flag for node1, indicating that you recognize it is dead and will not be able to reboot for some time.

Answer: C

QUESTION NO: 92

When configuring IPMP groups for use in the Sun Cluster 3.2 environment, which two statements are true? (Choose two.)

A. All interfaces must be in IPMP groups.
B. Each interface can be a member of multiple IPMP groups.
C. Adapters within the same IPMP group must be of the same transport type.
D. Adapters on different subnets must be in different IPMP groups.

Answer: C,D

QUESTION NO: 93

Which component of Sun Cluster 3.2 provides global naming consistency for shared storage?

A. the CCR
B. the RGM
C. DID devices
D. PxFS

Answer: C

QUESTION NO: 94

You have a three node cluster with three different resource groups. Rg-1 can only run on node1 and node3, Rg-2 can only run on node2 and node3, while Rg-3 can only run on node3. If Rg-1 or Rg-2 move to node3, Rg-3 has to be stopped.

How would you accomplish this ?

A. clrs set -p RG_affinities=--Rg-1,--Rg-2 Rg-3
B. clresourcegroup set -p RG_affinities=--Rg-3 +
C. clresourcegroup set -c RG_affinities=--Rg-1,Rg-2 Rg-3
D. clresourcegroup set -p RG_affinities=--Rg-1,--Rg-2 Rg-3

Answer: D

QUESTION NO: 95

The quorum mechanism and voting system are important parts of Sun Cluster. What is the majority formula to have a working cluster?

A. All the votes must be present.
B. At least 2 votes are required to form a cluster.
C. There must be a majority (more than 50 percent of all possible votes) to form a cluster or remain in a cluster.
D. Sun Cluster 3.2 automatically recalculates votes depending on the current configuration, so there is no need to worry about it.

Answer: C

QUESTION NO: 96

During the scinstall process the cluster administrator has to configure the cluster transport.

Which step does the cluster administrator have to perform in a typical cluster configuration?

A. Select the transport adapters.
B. Select the IP range for the private networks.
C. Construct the public network VLANS for a tagged VLAN.
D. Select the netmask for the cluster interconnect on a typical install.

Answer: A

QUESTION NO: 97

Which three are characteristics of the High Availability environment provided by Sun Cluster?
(Choose three.)

A. ability to run a single instance of Solaris across multiple nodes
B. automated application restart in the event of node failure
C. uninterrupted service in the event of any hardware failure
D. load sharing by scaling an application across several nodes
E. minimal service outage during hardware and software upgrades
F. possibility to combine different platform architectures in one system

Answer: B,D,E

QUESTION NO: 98

You have two local (non-shared) disks on a cluster node. You put one local metadb replica on one
disk and two local metadb replicas on the other disk.

What will happen if later the disk with the two replicas fails?

A. You are guaranteed that the node stays up and running, and can reboot. As long as Solaris
Volume Manager can find one valid copy of the configuration, the node will stay up.
B. You are guaranteed to stay up and running. However, if you reboot, you may have to manually
delete the broken metadb replicas before being able to join the cluster.
C. When Solaris Volume Manager discovers you have less than 50% of your local metadbs
remaining, it will fail immediately.
D. When Solaris Volume Manager discovers you have less than 50% of your local metadbs
remaining, it will prompt you to fix the broken ones and you can stay operational without rebooting.

Answer: C

QUESTION NO: 99

When would you use a Veritas Volume Manager CDS disk in the bootdg group?

A. All the time, it is the default.
B. You cannot use CDS disks in the bootdg group.
C. When you want all nodes to boot from the same device.
D. When you want to share your boot device between clusters.

Answer: B

QUESTION NO: 100

What is the maximum number of private interconnects supported in Sun Cluster 3.2?

A. Six
B. Four
C. One
D. Two

Answer: A

QUESTION NO: 101

A system administrator is preparing to install a new cluster.

What must be consulted to determine the most recent information about required patches?

A. http://www.sun.com/bigadmin/
B. http://sunsolve.sun.com/
C. The Sun Cluster 3.2 Administration guide on the distribution medium
D. http://www.sun.com/cluster/

Answer: B

QUESTION NO: 102

After consolidating multiple clusters into a single cluster you will need to run multiple HA Oracle instances in the same cluster. Which of the steps is unnecessary in an environment where the Oracle instances have to be accessed over a network?

A. Configure listener resources for the Oracle instances.
B. Configure the server resources for the Oracle instances.
C. Configure your global/failover file system in HAStoragePlus resources.
D. Configure the resources for Oracle Grid Control to administer the Oracle instances.

Answer: D

QUESTION NO: 103

Your customer has asked you to create a Veritas Volume Manager volume. They want the volume to be RAID level 1+0. The name of the group is oradg and they want the name of the volume to be oravol, and 3 Gbytes in size.

Which three statements are correct? (Choose three.)

A. RAID 1+0 requires at least 4 disks
B. The volume oravol will be a layered volume
C. vxmake -g oradg oravol layout=mirror
D. vxassist -g oradg make oravol 3g layout=stripe-mirror
E. vxassist -s oradg make oravol 3g layout=mirror-stripe

Answer: A,B,D

QUESTION NO: 104

Which two are prerequisites for installing Sun Cluster 3.2 software with the Java Enterprise System installer? (Choose two.)

A. All Java ES auxilliary software must be installed.
B. Solaris software must be installed and correctly patched.

C. The system boot disks must be correctly partitioned for the volume manager to be used.
D. A graphical interface display must be available on the node where installation is being done.

Answer: B,C

QUESTION NO: 105

The clrg show command shows the following output for pippo-rg resource group.

Resource group nfs-rg is configured to run only on node fachiro.

Resource Group: pippo-rg

RG_description:

RG_mode: Scalable

RG_state: Unmanaged

RG_project_name: default

RG_affinities: --nfs-rg

RG_SLM_type: manual

Auto_start_on_new_cluster: True

Fallback: False

Nodelist: fachiro zippo

Which is correct about the resource group configuration?

A. It is a scalable resource group that will not run on node fachiro if nfs-rg is online.
B. The nfs-rg resource group will stop when pippo-rg switch to node fachiro.
C. RG_affinities cannot be configured into a scalable resource group.
D. The clrg show + command would provide different information about the resource group.

Answer: A

QUESTION NO: 106

You have 12 ZFS file systems within the same ZFS zpool. You need them all to fail over together from node to node along with your application.

How should you configure your HAStoragePlus resource or resources to control this storage?

A. Configure 12 separate HAStoragePlus resources, one for each mount point.
B. Configure a single HAStoragePlus resource, with the Zpools property set to the name of the pool.
C. Configure 12 separate HAStoragePlus resources, each with the Zpools property set to the name of the same pool.
D. Configure a single HAStoragePlus file system with the FilesystemMountPoints property set to the list of 12 mount points.

Answer: B

QUESTION NO: 107

You want to give a non-root user the ability to online, enable, disable, and switch over elements of the cluster, but NOT to modify the configuration.

In which two ways can you accomplish this? (Choose two.)

A. Give that specific user the Solaris.cluster.admin authorization.
B. Give that user the ability to su to a non-root role that has the Solaris.cluster.admin authorization.
C. Give that user the ability to su to a non-root role that has the Solaris.cluster.modify authorization.
D. You do not have to do anything. Every user can accomplish what you want since they will naturally all inherit the Solaris.cluster.read authorization from the Basic Solaris User profile.

Answer: A,B

QUESTION NO: 108

An Oracle DBA wishes to shut down a database, which is under Sun Cluster 3.2 control to perform maintenance. Which two procedures could the cluster administrator use to enable the DBA to perform this task? (Choose two)

A. clrs unmonitor oracle-setver-res

(DBA shuts down the database and performs maintenance)

clrs monitor oracle-server-res

B. clrg suspend oracle-rg

(DBA shuts down the database and performs maintenance)

clrg resume oracle-rg

C. clrg switch -n node2 oracle-rg

(DBA shuts down the database and performs maintenance)

clrg switch -n node1 oracle-rg

D. sqlplus / as sysdba

shutdown immediate

(DBA performs maintenance)

sqlplus / as sysdba

startup

E. clrs disable oracle-lsnr-res

(DBA shuts down the listener and performs maintenance)

clrs disable oracle-lsnr-res

Answer: A,B

QUESTION NO: 109

Your customer has just created a new Veritas Volume Manager volume named oravol in the already registered disk group oradg. Now they are trying to create a UFS filesystem on top of the volume and keep getting the following error message after executing the command:

newfs /dev/vx/rdsk/oradg/oravol

/dev/vx/rdsk/oradg/oravol: No such device or address

Which two statements describe the cause of the error? (Choose two.)

A. You did not run cldg sync oradg.

B. You typed the newfs command incorrectly.

C. The newfs command must be run on one of the other nodes.

D. The global namespace does NOT know about this new volume on the other nodes.

Answer: A,D

QUESTION NO: 110

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

Trying 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 **69,000** Satisfied Customers Worldwide

★ Multi-Platform capabilities - **Windows, Mac, Android, iPhone, iPod, iPad, Kindle**

Need Help

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

One Year Free Update
Free update is available within One Year after your purchase. After One Year, you will get 50% discounts for updating. And we are proud to boast a 24/7 efficient Customer Support system via Email.

Money Back Guarantee
To ensure that you are spending on quality products, we provide 100% money back guarantee for 30 days from the date of purchase.

Security & Privacy
We respect customer privacy. We use McAfee's security service to provide you with utmost security for your personal information & peace of mind.

Guarantee & Policy | Privacy & Policy | Terms & Conditions