

Oracle

Exam 1z0-820

Upgrade to Oracle Solaris 11 System Administrator

Version: Demo

[Total Questions: 10]

Question No : 1

View the Exhibit and review the zpool and ZFS configuration information from your system.

```

pool: pool1
state: ONLINE
scan: none requested
config:

    NAME                STATE                READ WRITE CKSUM
    pool1                ONLINE              0   0   0
      mirror-0          ONLINE              0   0   0
        c4t0d0          ONLINE              -   -   -
        c4t1d0          ONLINE              -   -   -
      mirror-1          ONLINE              0   0   0
        c4t2d0          ONLINE              -   -   -
        c4t3d0          ONLINE              -   -   -

errors: No known data errors

pool: rpool
state: ONLINE
scan: none requested
config:

    NAME                STATE                READ WRITE CKSUM
    rpool                ONLINE              0   0   0
      c3t0d0s0          ONLINE              0   0   0

errors: No known data errors
NAME                                USED    AVAIL    REFER    MOUNTPOINT
pool1                                138K    7.81G    32K      /pool1
pool1/prod_data                       31K    7.81G    31K      /prod_data
rpool                                11.6G   4.04G    34.5K    /rpool
rpool/ROOT                            9.94G   4.04G    31K      legacy
rpool/ROOT/solaris                    9.94G   4.04G    9.70G    /
rpool/dump                             630M    4.05G    611M    -
rpool/export                          6.07M   4.04G    32K      /export
rpool/export/home                     6.04M   4.04G    32K      /export/home

```

Identify the correct procedure for breaking the /prod_data mirror, removing c4t0d0 and c4t2d0, and making the data on c4t0d0 and c4t2d0 accessible under the dev_data mount point.

- A.** zpool split pool1 pool2 c4t0d0 c4t2d0
zpool import pool2
zfs set mountpoint = /dev_data pool2/prod_data
- B.** zpool detach pool1 pool2
zpool attach pool2
zfs set mountpoint=/dev_data pool2/prod_data
- C.** zpool split pool1/prod_data -n pool2/dev_data
zfs set mountpoint = /dev_data pool2/prod_data
- D.** zpool split pool1 pool2 c4t0d0 c4t2d0
zpool import pool2

Answer: D

Explanation: In this Solaris release, you can use the zpool split command to split a mirrored storage pool, which detaches a disk or disks in the original mirrored pool to create

another identical pool.

After the split operation, import the new pool.

Reference: Oracle Solaris ZFS Administration Guide, Creating a New Pool By Splitting a Mirrored ZFS Storage Pool

Question No : 2

Which three are true of a solaris0 branded zone running on an Oracle Solaris 11 system? (Choose three.)

- A. if there are any native non-global zones on the system that you want to migrate, these zones will be migrated into non-global zones on the Solaris 11 system
- B. The non-global zone must reside on its own ZFS dataset; only ZFS is supported
- C. The non-global zone can reside on a ZFS dataset or a UFS file system
- D. The solaris0 branded zone can be either a sparse or a whole root non-global zone model.
- E. The default is the exclusive-IP type.
- F. When migrated from a stand-alone Solaris system to a solaris0 branded zone, the host ID changes to the host ID of the new machine The zone cannot be configured to use the host ID of the original system
- G. A sparse root zone cannot be converted into a solaris0 branded zone.
- H. The solaris0 branded zone cannot be an NFS server

Answer: B,E,H

Question No : 3

You want the system to generate core files when an application crashes, but you want only root to be able to view the core files. The current core configuration is:

global core file pattern: /var/core/core %f.%p

global core file content: default

init core file pattern: /var/core/pprocess/core.%f.%pp

init core file content: default

global core dumps: enabled

per—process core dumps: enabled

global setid core dumps disabled

per-process setid core dumps: disabled

global core dump logging disabled

Select the option that describes the change that needs to be made to the core dump configuration to meet this requirement.

- A. coreadm —d process
- B. coreadm -i /var/core/core.%f %p
- C. coreadm—d global
- D. coreadm -d process -d global -e global-setid -e proc-setid
- E. Make this change to the /etc/coreadm.conf file: per-process core dumps: disabled

Answer: A

Question No : 4

United States of America export laws include restrictions on cryptography.

Identify the two methods with which these restrictions are accommodated in the Oracle Solaris 11 Cryptographic Framework.

- A. Corporations must utilize signed X.509 v3 certificates.
- B. A third-party provider object must be signed with a certificate issued by Oracle.
- C. Loadable kernel software modules must register using the Cryptographic Framework SPI.
- D. Third-party providers must utilize X.509 v3 certificates signed by trusted Root Certification Authorities.
- E. Systems destined for embargoed countries utilize loadable kernel software modules that restrict encryption to 64 bit keys.

Answer: B,C

Explanation:

B: Binary Signatures for Third-Party Software

The elfsign command provides a means to sign providers to be used with the Oracle Solaris Cryptographic Framework. Typically, this command is run by the developer of a

provider.

The `elfsign` command has subcommands to request a certificate from Sun and to sign binaries. Another subcommand verifies the signature. Unsigned binaries cannot be used by the Oracle Solaris Cryptographic Framework. To sign one or more providers requires the certificate from Sun and the private key that was used to request the certificate.

C: Export law in the United States requires that the use of open cryptographic interfaces be restricted. The Oracle Solaris Cryptographic Framework satisfies the current law by requiring that kernel cryptographic providers and PKCS #11 cryptographic providers be signed.

Reference: System Administration Guide: Security Services, Oracle Solaris Cryptographic Framework

Question No : 5

The current ZFS configuration on serverA is:

Pool1

c3t2d0

c3t3d0

Pool2

c3t4d0

c3t5d0

The ZFS configuration on serverB is:

Pool1

mirror-0

c3t2d0

c3t3d0

mirror-1

c3t4d0

c3t5d0

You need to change serverA's ZFS configuration to match serverB's configuration.

Which option should you choose to modify the configuration on serverA so that it matches serverB's configuration?

- A. zpool destroy pool2
zpool attach pool1 c3t4d0 ct3t5d0
- B. zpool destroy pool2
zpool attach pool1 c3t2d0 c3t2d0 c3t4d0 c3t5d
- C. zpool destroy pool2
zpool add pool1 c3t4d0 c3t5d0
- D. zpool destroy pool2
zpool mirror pool1 pool2
- E. zpool destroy pool2
zpool attach pool1 c3t2d0 c3t4d0
zpool attach pool1 c3t3d0 c3t5d0
- F. zpool destroy pool1
zpool destroy pool2
zpool create pool1 mirror c3t2d0 c3t3d0 c3t4d0 c3t5d0

Answer: F

Explanation: We destroy both pools. Then we recreate the mirror.

Note:

```
zpool attach [-f] pool device new_device
```

Attaches new_device to an existing zpool device. The existing device cannot be part of a raidz configuration. If device is not currently part of a mirrored configuration, device automatically transforms into a two-way mirror of device and new_device. If device is part of a two-way mirror, attaching new_device creates a three-way mirror, and so on. In either case, new_device begins to resilver immediately.

Reference: man zpool

Which modification needs to be made to the Service Management Facility before you publish a new package to the IPS repository?

- A. The pkg.depotd must be disabled.
- B. The pkg/readonly property for the application/pkg/server service must be set to false
- C. The pkg/writable_root property for the application/pkg/server service must be set to true.
- D. The pkg/image_root property for the application/pkg/server service must be set to the location of the repository.

Answer: D

Explanation: pkg/image_root

(astring) The path to the image whose file information will be used as a cache for file data.

Reference: man pkg.depotd

Question No : 7

The COMSTAR framework provides support for the iSCSI protocol.

Identify three true statements describing the COMSTAR framework. (Choose three.)

- A. One IP port can handle multiple iSCSI target devices.
- B. The disk volume provided by the server is referred to as the LUN and can be accessed by an iSCSI initiator.
- C. The disk volume provided by the server is referred to as the target.
- D. The iSCSI initiator is configured on the server.
- E The iSCSI initiator is configured on the client.

Answer: A,C,D

Question No : 8

Select the two statements that correctly describe the operation of NWAM.

- A. If a location is explicitly enabled, it remains active until explicitly changed.
- B. Wireless security keys can be configured by using the nwammgr command.
- C. NWAM stores profile information in /etc/ipadm/ipadm.conf and /etc/dladm/datalink.conf.
- D. Multiple locations may be automatically activated in systems with multiple network

interface cards.

E. Interface NCU Properties "float" and are automatically attached to the highest priority Link NCU Property.

F. If the DefaultFixed NCP is enabled, persistent configuration, stored in /etc/ipadm.conf and /etc/dladm/datalink.conf is used.

Answer: A,D

Explanation: A: Conditional and system locations can be manually activated, which means that the location remains active until explicitly disabled.

D: A location comprises certain elements of a network configuration, for example a name service

and firewall settings, that are applied together, when required. You can create multiple locations

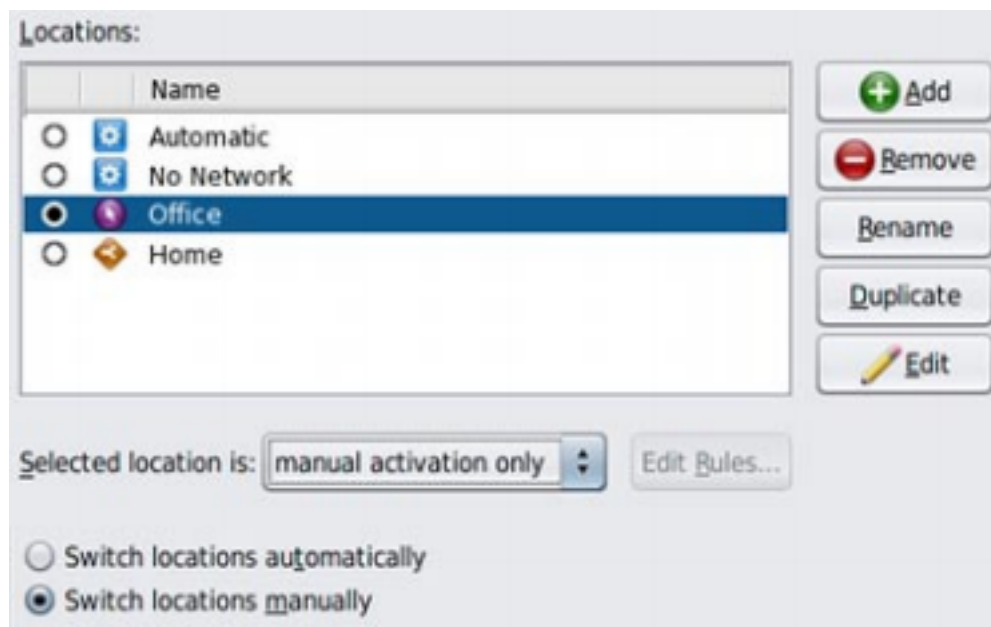
for various uses. For example, one location can be used when you are connected at the office by

using the company intranet. Another location can be used at home when you are connected to

the public Internet by using a wireless access point. Locations can be activated manually or automatically, according to environmental conditions, such as the IP address that is

obtained

by a network connection.



Reference: Oracle Solaris Administration: Network Interfaces and NetworkVirtualization, Activating and Deactivating Profiles

Reference: Oracle Solaris Administration: Network Interfaces and NetworkVirtualization, Creating and Managing Locations

Question No : 9

You are using Wireshark or snoop to observe network traffic, by using the observability node /dev/net/bridgeA. The bridge was created with: `dladm create-bridge -1 net1 -1 net2 bridgeA`

The captured packets are those_____.

- A. sent, prior to VLAN tag modifications
- B. received, after VLAN tag modifications
- C. received, prior to VLAN tag modifications
- D. sent and received, after VLAN tag modifications

Answer: A

Question No : 10

Your mentor suggests using the `dladm rename-link` command to rename the network datalinks.

What are the two advantages of following this advice?

- A. It can clarify which network interface has what purpose.
- B. It can simplify specifying the network interface with the `dladm modify-aggr` command.
- C. It can simplify specifying the network interface with the `dladm modify-bridge` command.
- D. It can simplify IP filter rule changes if the network interface is replaced with a different type.
- E. It can prevent accidental deletion of the network interface with the `dladm delete-phys` command.
- F. It can prevent accidental deletion of the network interface configuration with the `ipadm delete-addr` command.

Answer: A,E

Explanation: A: To rename the `bge0` link to `mgmt0`, enter the following command:

```
# dladm rename-link bge0 mgmt0
```

E: Consider that the `bge0` device, whose link was named `mgmt0` as shown in the previous

example, needs to be replaced with a ce0 device because of a hardware failure. The bge0 NIC is physically removed, and replaced with a new ce0 NIC. To associate the newly added ce0 device with the mgmt0 configuration previously associated with bge0, enter the following command:

```
# dladm rename-link ce0 mgmt0
```

Note: How to Rename a Datalink

Use this procedure if you want to change a datalink name to a customized name. For example, some of the datalinks in upgraded system might have retained legacy hardware-based names and you want to change these names to generic ones.

Note: `dladm rename-link [-R root-dir] link new-link`

Rename link to new-link. This is used to give a link a meaningful name, or to associate existing link configuration such as link properties of a removed device with a new device.