



Oracle

Exam 1z0-820

Upgrade to Oracle Solaris 11 System Administrator

Version: 7.0

[Total Questions: 133]

Question No : 1

You are creating a non-global zone on your system.

Which option assigns a zpool to a non-global zone, and gives the zone administrator permission to create zfs file system in that zpool?

A. While creating the non-global zone, make the following entry:

```
add device
set match=/dev/rdisk/c4t0d0
end
```

Boot the zone and log in the zone as root. Create the zpool:

```
zpool create pool2 c4t0d0
```

In the non-global zone, root can now create ZFS file system in the pool2 zpool

B. In the global zone, create the zpool:

```
global# zpool create pool2 c4t1d0
```

While creating the no-global zone, make the following entry:

```
add dataset
set name=pool2
end
add fs
set dir=pool1
set special=pool1
set type=zfspool1
end
```

Boot the zone, log in the zone as root, and create the zfs file system in the pool2 zpool.

C. In the global zone, create the zpool:

```
global#zpool create pool2 c4t1d0
```

While creating the global zone, make the following entry:

```
add dataset
set name=pool2
end
```

Boot the zone, log in to the zone as root and create the zfs file systems in the pool2 zpool.

D. In the global zone, create the zpool and the ZFS file systems that you want to use in the non-global zone:

```
global#zpool create pool2 c4t1d0
global#zfs create pool2/data
```

While creating the non-global zone, make the following entry for each ZFS file system that you want to make available in the zone:

```
add fs
set dir=/data
set special=pool2/data
set type=zfs
end
```

E. Create the zpool in the global zone:

```
global#zpool create pool2 c4t1d0
```

Boot the non-global zone, log in to the zone as root, and issue this command to delegate ZFS permissions to root:

```
non-global zone# zfs allow root create , destroy, mount pool2
```

Log in to the non-global zone create ZFS file systems in the pool2 zpool.

Answer: C

Reference: Oracle Solaris ZFS Administration Guide, Adding ZFS File Systems to a Non-Global Zone

Question No : 2

You are using Wireshark or snoop to observe network traffic, by using the observability node /dev/net/bndgeA. 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 : 3

You are going to use the- Automated installer (AI) to install a non global zone named zone1. You have created a custom manifest for the non-global zone and named it zone1manifest

Which command will you use to add this custom manifest to the s11-sparc install service and associate this custom manifest with the non-global zone?

- A. `installadm create-profile -n s11-sparc -f /tmp/zone1manifest.xml - c`
- B. `installadm create-manifest -n s11-sparc -f /tmp/zone1manifest.xml -m`

C. `installadm create-client -n s11-sparc -f /tmp/zone1manifest.xml -m zone1manifest -c zonename="zone1"`

D. `installadm create-service -n s11-sparc -f /tmp/zone1manifest.xml -m zone1manifest -c zonename="zone1"`

Answer: B

Explanation: `installadm add-manifest`

Associates manifests with a specific install service, thus making the manifests available on the network, independently from creating a service. When publishing a non-default manifest, it is required to associate criteria either via criteria entered on the command line (-c) or via a criteria XML file (-C).

Reference: `installadm`

Question No : 4

The crash dump notification on your server is:

```
Dump content: kernel and current process pages
Dump device: /dev/zvol/dsk/rpool/dump (dedicated)
Savecore directory: /var/crash
Savecore enabled: no
Save compressed: on

The files in the /var/crash directory are:
bounds      vmdump.0
```

Documentation states that there should be two core files for each crash dump in the `/var/crash` directory named `vmdump.0`

Which command should you choose to display these two files?

- A. `savecore -f vmdump.0`
- B. `dumpadm uncompressed`
- C. `gunzip vmdump.0`
- D. `dumpadm -z off`

Answer: A

Explanation: Decompress using savecore -f vmdump.0

savecore - save a crash dump of the operating system

-f **dumpfile** Attempt to save a crash dump from the specified file instead of from the system's current dump device. This option may be useful if the information stored on the dump device has been copied to an on-disk file by means of the dd(1M) command.

Question No : 5

You want to deploy Oracle Solaris 11 with the Automated Installer (AI). You need to make sure that your server and network meet the requirements for using AI. Choose the three options that describe the requirements for using AI. (Choose three.)

- A. You must set up DHCP. The DHCP server and AI install server can be the same machine or two different machines.
- B. You can create only one manifest per install service. If you need more than one manifest, create multiple install services.
- C. The minimum requirement to operate as an AI install server is 1 GB of memory.
- D. If two client machines have different architectures and need to be installed with the same version of the Oracle Solaris 11 OS(then create two AI manifests and a single install service.
- E. Load the install/installadm package on the AI server.

Answer: A,D,E

Question No : 6

Review the zpool information on your server:

NAME	SIZE	ALLOC	FREE	CAP	DEDUP	HEALTH	ALTROOT
pool1	3.97G	141K	3.97G	0%	1.00x	ONLINE	-
rpool	15.9G	11.6G	4.31G	72%	1.00x	ONLINE	-

You enabled the deduplication properly on the poolH/data file system, yet the information on the pool shows a DEDUP RATIO of 1.00x. Which two options describe why the data duplication ratio on poolM still shows 1.00x and not a higher value? (Choose two.)

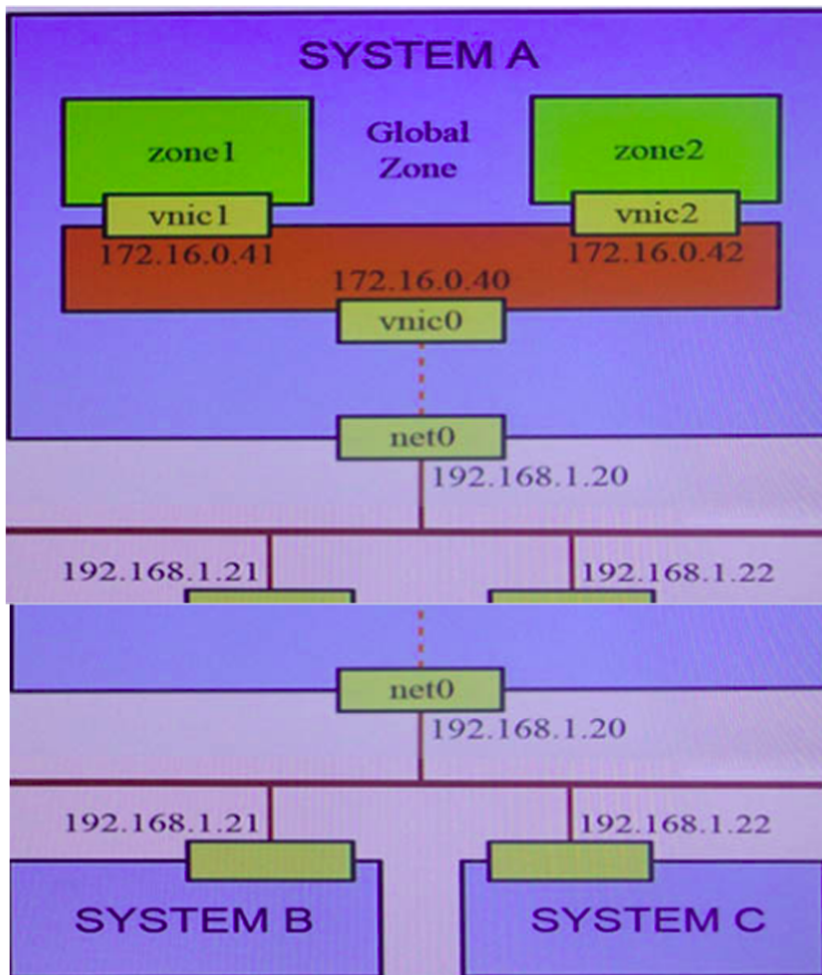
- A. Most of the data in the poolM/data file system is not duplicate data.
- B. You have the copies property set to a value of greater than 1 on your ZFS file system.
- C. You have set the dedup property on the ZFS file system, but have not set the dedupratio property on the zpool.
- D. You have set the deduplication property on the ZFS file system, but have not set the dedupratio property on the zpool.
- E. You have set the dedup property on the ZFS file system, but have not added any data to the file system since the dedup property was set.

Answer: B,C

Question No : 7

On SystemA, you are troubleshooting a new installation of a virtual network. Local zones cannot connect with external systems.

View the Exhibit.



Logged in to zone1 or system vty you can successfully ping 172.16.0.40 and 172.16.0.41, but there is no response from any systems on the 192.168.1.x network.

Similarly, logged in to zone2 of SystemA, you can successfully ping 172.16.0.40 and 172.16.0.41, but there is no response from any systems on the 192.168.1.x network.

However, logged in to the global zone of SystemA, you can successfully ping 172.16.0.41 and 172.16.0.42, as well as 192.168.1.21 and 192.168.1.22.

Identify two possible explanations. (Choose two.)

- A. forwarding=on has not been configured for SystemA and SystemB.
- B. forwarding=on has not been configured for the global zone of SystemA.
- C. There are missing entries for SystemA and SystemB in /etc/inet/hosts.
- D. forwarding=on has not been configured for zone1 and zone2 of SystemA.
- E. ARP table entries on SystemA and SystemB are incomplete for vnic1 and vnic2.
- F. SystemB and SystemC require a static route to the 172.16.0.x network via 192.168.1.20.

Answer: B,E

Question No : 8

Consider the following rule file for use with the Basic Audit Reporting Tool (BART).

CHECK all

IGNORE dirmtime

/etc/security

/etc/notices

IGNORE contents

/export/home

IGNORE mtime size contents

/var

CHECK

You are using BART to detect inappropriate changes to the file system.

Identify the two correct statements describing the attributes recorded.

- A. /var/dhcp Attribute: size uid gid mode acl
- B. /etc/hosts Attributes: size uid gid mode acl intime dest
- C. /var/spool/mqueue Attribute: size uid gid mode acl dirmtime
- D. /etc/security/exec_attr Attribute: size uid mode acl mtime devnode
- E. /export/home/kate/.profile Attributes: uid gid mode acl dirmtime
- F. /export/home/rick/.profile Attributes: size uid gid mode acl mtime contents

Answer: D,F

Explanation: D: According to line /etc/security

F: According to line /export/home

Not E: According to line IGNORE dirmtime

Note: In default mode, the bart compare command, as shown in the following example, checks all the files installed on the system, with the exception of modified directory

timestamps (dirmtime):

CHECK all

IGNORE dirmtime

Note 2: The Basic Audit Reporting Tool (BART) feature of Oracle Solaris enables you to comprehensively validate systems by performing file-level checks of a system over time. By creating BART manifests, you can easily and reliably gather information about the components of the software stack that is installed on deployed systems.

BART is a useful tool for integrity management on one system or on a network of systems.

Reference: Oracle Solaris Administration: Security Services, BART Manifests, Rules Files, and Reports (Reference)

Question No : 9

Which two accurately describe the Image Packaging System (IPS) components? (Choose two.)

- A. Every package is represented by an FMRI, consisting of a publisher, a name, and a version, with the scheme.
- B. Every package is represented by a manifest, consisting of a publisher, a name, and a version, with the scheme.
- C. A publisher is an entity that develops and constructs packages.
- D. a repository is an entity that publishes and provides access to packages.

Answer: A,C

Question No : 10

You are setting up a local IPS package repository on your Oracle Solaris 11 server; solans example.com

You want to point the existing local IPS publisher to the new local IPS repository located in /repo. These are the steps that you have followed:

1 Download and rsync the contents of the Oracle Solaris 11 repository ISO image to the /repo directory.

2. Configure the repository server service properties. The svcprop command displays the IPS-related properties:

```
pkg/inst_root astng /repo pkg/readonly boolean true
```

The ls command displays the contents of the /repo directory:

```
#ls /repo
```

```
pkg5 repository publisher
```

The svcs command shows the svc:/application/pkg/server:default service is online.

The pkg publisher command still displays:

```
PUBLISHERTYPESTATUSURI solarisoriQinonlinehttp://DkQ.Qracle.conVsolaris/release/
```

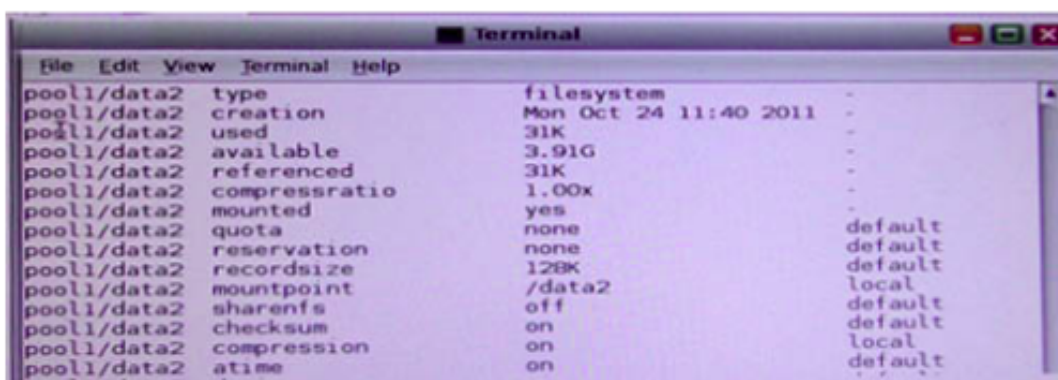
Which step needs to be performed to set the local IPS publisher to the local IPS repository /repo?

- A. issue the pkgrepo refresh -s command to refresh the repository.
- B. Restart the svc:/application/pkg/server:default service.
- C. issue the pkg set-publisher command to set the new repository location.
- D. issue the pkgrepo rebuild command to rebuild the repository.
- E. issue the pkgrepo set command to set the new repository location-

Answer: E

Question No : 11

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



The application development team requested an up-to-date copy of the data from the /prod_data file system. You decide to give the team one of the disk drives containing the data by breaking the mirror, removing the disk c4t1d0 and mounting c4t1 d0 under a new mount point named /dev_data. Identify the correct procedure for breaking the mirror, removing c4t1d0. and making the data on that drive accessible under the /dev_data mount point

- A. zfs destroy pool /prod_data
zfs create pool1/prod_data c4t1d0 zfs create pool1/dev_data c4t1d0
- B. zfs split pool1/prod_data
zfs mount -F zfs pool1/dev_data c4t1d0
- C. zpool split pool1 pool2 zpool import pool2 zfs set mountpoint=/dev_data pool2/prod_data
- D. zpool split pool1/prod_data -n pool2/dev_data
zpool import -o mountpoint=/dev_data pool2/dev_data
zfs split pool1/prod_data -n pool2/dev_data zfs set mountpoint=/dev_data pool2/dev_data

Answer: C

Question No : 12

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 : 13

You display the IP Interface information with `ipmpstat -i`

Which two characteristics are indicated by characters that may be included in the `FLAGS` column?

- A. default route
- B. IP forwarding enabled
- C. allocated to global zone
- D. unusable due to being inactive
- E. nominated to send/receive IPv4 multicast for its IPMP group

Answer: D,E

Explanation: The `ipmpstat` command concisely displays information about the IPMP subsystem. It supports five different output modes, each of which provides a different view of the IPMP subsystem (address, group, interface, probe, and target), described below.

-i

Display IP interface information (“interface” output mode).

Interface Mode

Interface mode displays the state of all IP interfaces that are tracked by `in.mpathd` on the system. The following output field is one of the supported:

FLAGS

Assorted information about the IP interface:

i

(D) Unusable due to being INACTIVE.

s

Marked STANDBY.

m

(E) Nominated to send/receive IPv4 multicast for its IPMP group.

b

Nominated to send/receive IPv4 broadcast for its IPMP group.

M

Nominated to send/receive IPv6 multicast for its IPMP group.

d

Unusable due to being down.

h

Unusable due to being brought OFFLINE by in.mpathd because of a duplicate hardware address.

Reference: man ipmpstat

Question No : 14

When setting up Automated Installer (AI) clients, an interactive tool can be used to generate a custom system configuration profile. The profile will specify the time zone, date and time, user and root accounts, and name services used for an AI client installation. This interactive tool will prompt you to enter the client information and an SC profile (XML file) will be created.

Which interactive tool can be used to generate this custom configuration?

- A. sys-unconfig
- B. installadm set-criteria
- C. sysconfig create-profile
- D. installadm create-profile

Answer: C

Question No : 15

A non-global zone named testzone is currently running.

Which option would you choose to dynamically set the CPU shares for the zone to two shares?

A. While logged in to the global zone, enter:

`prctl -n zone.cpu-shares -v 2 -r -i zone testzone`

B. While logged in to the global zone, enter: `zonecfg -z testzone add rctl set name=zone cpu-shares`

`setvalue=(priv=privilegedJimit=2,action=none)`

C. While logged in to the global zone, enter: `prctl -n 2 zone cpu-shares -i zone testzone`

D. While logged in to the global zone, enter: `zonecfg -z testzone add rctl set name-zone cpu-shares set value=2`

E. While logged in to testzone, enter:

`prctl -n zone.cpu-shares -v 2 -r -i zone testzone`

Answer: D

Question No : 16

When you issue the `gzip` command, the "gzip: command not found" message is displayed. You need to install the `gzip` utility on your system. Which command would you use to check if the `gzip` utility is available from the default publisher for installation?

A. `pkg info|grep gzip`

B. `pkg list SUNWgzip`

C. `pkg contents gzip`

D. `pkg search gzip`

Answer: D

Question No : 17

You execute the command:

`usermod -K limitpvn=all,\!file_wnte guest`



What is the result of this command?

- A. The guest account cannot write any files.
- B. The guest account can assume any role except the file_write role.
- C. Starting at next login, the guest account will be unable to write any files.
- D. The guest account cannot assume a role that includes file_write privileges.
- E. Unless the guest account assumes the limitpriv role, it cannot write any files
- F. An error message is displayed, indicating that "file_wnte" is not a valid execution attribute

Answer: C

Question No : 18

A change in your company's security policy now requires an audit trail of all administrators assuming the sysadm role, capturing:

-  Executed commands, including options
-  Logins and logouts


There are two command necessary to accomplish this change. One is a rolemod command. What is the other?

- A. auditconfig set policy=argv
- B. auditconfig -setpolicy +argv
- C. auditconfig -setflags lo, ex sysadm
- D. auditconfig set flags=lo, ex sysadm

Answer: B

Explanation: Audit Significant Events in Addition to Login/Logout (see step 2 below)


Use this procedure to audit administrative commands, attempts to invade the system, and other significant events as specified by your site security policy.

-  Audit all uses of privileged commands by users and roles.

For all users and roles, add the AUE_PFEXEC audit event to their preselection mask.

```
# usermod -K audit_flags=lo,ps:no username
```

```
# rolemod -K audit_flags=lo,ps:no rolename
```

-  Record the arguments to audited commands.

```
# auditconfig -setpolicy +argv
```

3- Record the environment in which audited commands are executed.

```
# auditconfig -setpolicy +arge
```

Note: [-t] -setpolicy [+|-]policy_flag[,policy_flag ...]

Set the kernel audit policy. A policy policy_flag is literal strings that denotes an audit policy. A prefix of + adds the policies specified to the current audit policies. A prefix of - removes

the policies specified from the current audit policies. No policies can be set from a local zone unless the perzone policy is first set from the global zone.

Reference: Oracle Solaris 11 Security Guidelines, Audit Significant Events in Addition to Login/Logout

Question No : 19

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 : 20

Select the five tasks that need to be performed on the Automated Installer (AI) install server before setting up the client.

- A. Create a local IPS repository on the AI Install server and start the repository server service, the publisher origin to the repository file.
- B. Set up a IP address on the AI install server.
- C. The DHCP server must be enabled on the install server and must provide the DHCP service for the clients.
- D. DHCP must be available on the network for the Install server and the clients, but the install server does not need to be the DHCP server.
- E. Download the AI boot image. The image must be the same version as the Oracle Solaris

OS that you plan to install on the client.

F. Download the text install image into the IPS repository.

G. Install the AI installation tools.

H. Create the AI install service. Specify the path to the AI network boot image ISO file and the path where the AI net image ISO file should be unpacked.

I. Create the AI install service. Specify the path to the AI network boot image ISO file and the path to the IPS repository.

Answer: B,D,F,G,I

Explanation: B: Configure the AI install server to use a static IP address and default route.

D: The create-service command can set up DHCP on the AI install server. If you want to set up a separate DHCP server or configure an existing DHCP server for use with AI. The DHCP server must be able to provide DNS information to the systems to be installed.

E: An automated installation of a client over the network consists of the following high-level steps:

1. The client system boots over the network and gets its network configuration and the location of the install server from the DHCP server.

2. The install server provides a boot image to the client.

3. Characteristics of the client determine which installation instructions and which system configuration instructions are used to install the client.

4. The Oracle Solaris 11 OS is installed on the client, pulling packages from the package repository specified by the installation instructions in the AI install service.

G: Install the AI tool set.

Use the installadm create-service command to create an AI install service. Give the service a meaningful name, and specify the path where you want the service created. Specify the source of the network boot image (net image) package or ISO file.

```
installadm create-service [-n svcname]
```

```
[-s FMRI_or_ISO] [-d imagepath]
```

```
-d imagepath
```

The imagepath is the location of the new install service. The install-image/solaris-auto-install package is installed to this location, or the specified ISO file is expanded at this location.

Reference: Installing Oracle Solaris 11 Systems, Create an AI Install Service

Question No : 21

This icon is displayed on the desktop of a laptop computer, which is running Oracle Solaris 11.



Which two statements describe the Information conveyed by this Icon?

- A. NWAM is disabled.
- B. NWAM is automatically configuring the network.
- C. The wireless network card is manually configured and operational.
- D. The wireless network card is manually configured but not operational.
- E. The wireless network card is automatically configured and operational.
- F. The wireless network card is automatically configured but not operational.

Answer: B,C

Explanation: B: The Network Status notification icon is only displayed on the desktop if you are using NWAM to automatically configure your network.

C: All online (Wireless)

Indicates all manually enabled connections in the enabled network profile are online and that the required number of connections in the enabled profile group (if such a group exists) are online. The required number is the same as those described for the All online (Wired) status.

Note that at least one online connection is wireless.

Reference: Oracle Solaris Administration: Network Interfaces and Network Virtualization ,

Checking the Status of Your Network Connection

Question No : 22

View the Exhibit

```

<?xml version="1.0" encoding="UTF-8"?>
<!--
  Copyright (c) 2011, Oracle and/or its affiliates. All rights reserved.
-->
<!DOCTYPE auto_install SYSTEM "file:///usr/share/install/ai.dtd.1">
<auto_install>
  <ai_instance name="zone_default">
    <target>
      <logical>
        <zpool name="rpool">
          <filesystem name="export" mountpoint="/export"/>
          <filesystem name="export/home"/>
          <be name="solaris">
            <options>
              <option name="compression" value="on"/>
            </options>
          </be>
        </zpool>
      </logical>
    </target>

    <software type="IPS">
      <destination>
        <image>
          <!-- Specify locales to install -->
          <facet set="false">facet.locale.*</facet>
          <facet set="true">facet.locale.de</facet>
          <facet set="true">facet.locale.de_DE</facet>
          <facet set="true">facet.locale.en</facet>
          <facet set="true">facet.locale.en_US</facet>
          <facet set="true">facet.locale.es</facet>
          <facet set="true">facet.locale.es_ES</facet>
          <facet set="true">facet.locale.fr</facet>
          <facet set="true">facet.locale.fr_FR</facet>
          <facet set="true">facet.locale.it</facet>
          <facet set="true">facet.locale.it_IT</facet>
          <facet set="true">facet.locale.ja</facet>
          <facet set="true">facet.locale.ja_*</facet>
          <facet set="true">facet.locale.ko</facet>
          <facet set="true">facet.locale.ko_*</facet>
          <facet set="true">facet.locale.pt</facet>
          <facet set="true">facet.locale.pt_BR</facet>
          <facet set="true">facet.locale.zh</facet>
          <facet set="true">facet.locale.zh_CN</facet>
          <facet set="true">facet.locale.zh_TW</facet>
        </image>
      </destination>
      <software_data action="install">
        <name>pkg:/group/system/solaris-small-server</name>...
      </software_data>
    </software>
  </ai_instance>
</auto_install>

```

The file came from your Automated Installer (AI) Install server.

The file is_____.

- A. an AI SC profile for non-global zones
- B. the default AI config file for non-global zones
- C. the default AI manifest for non-global zones
- D. a custom AI manifest

Answer: C

Question No : 23

You are troubleshooting the failure of a computer to mount an NFS file system hosted by a server (hostname mars) in the local area network.

Select the three commands that will enable you to identify the problem.

- A. ping -s mars
- B. cat /etc/vfstab
- C. cat /etc/dfs/dfstab
- D. sharemgr show -v
- E. showmount -e mars
- F. rpcinfo -s mars | egrep 'nfs|mountd'

Answer: A,E,F

Explanation:

mount: mount-point /DS9 does not exist.

To solve the mount point error condition, check that the mount point exists on the client. Check the spelling of the mount point on the command line or in the /etc/vfstab file (B) on the client, or comment out the entry and reboot the system.

Note: The /etc/vfstab file lists all the file systems to be automatically mounted at system boot time, with the exception of the /etc/mnttab and /var/run file systems.

E: showmount

This command displays all clients that have remotely mounted file systems that are shared from an NFS server, or only the file systems that are mounted by clients, or the shared file systems with the client access information. The command syntax is:

```
showmount [ -ade ] [ hostname ]
```

where -a prints a list of all the remote mounts (each entry includes the client name and the directory), -d prints a list of the directories that are remotely mounted by clients, -e prints a list of the files shared (or exported), and hostname selects the NFS server to gather the information from. If hostname is not specified the local host is queried.

F: * mountd Daemon

This daemon handles file-system mount requests from remote systems and provides access control. The mountd daemon checks /etc/dfs/sharetab to determine which file systems are available for remote mounting and which systems are allowed to do the remote mounting.

* Commands for Troubleshooting NFS Problems

These commands can be useful when troubleshooting NFS problems.

rpcinfo Command

This command generates information about the RPC service that is running on a system.

Reference: man showmount

Question No : 24

You need a copy of the zone configuration installed on your server for disaster recovery documentation. Which two commands would display the current configuration of the zone named testzone? (Choose two)

- A. zonecfg -z testzone info
- B. zoneinfo -z testzone
- C. zoneadm -z testzone info
- D. zonecfg -z testzone export
- E. zonecfg testzone info
- F. cat /etc/zones/testzone.cfg

Answer: D,E

Question No : 25

Select the two correct descriptions of the configuration of a default installation of Oracle Solaris 11. (Choose two.)

- A. The root account no longer exists.
- B. Access to the root account is restricted by RBAC.
- C. The Solaris Management Console GUI (SMC) can be used to assign the root role to a user.
- D. telnet, ftp, and other services are disabled for network access, but root can log in by

using SSH.

E. Any logged in user, who knows the password for the root account, can modify the system configuration.

F. `usermod -R +root alice` will allow user alice to assume the role of root to modify the system configuration.

Answer: B,E

Question No : 26

You run the command `dlstat show-link -r`.

Select the two correct statements regarding the information displayed in the INTRS column.

A. No value is listed for virtual network interfaces.

B. A value of 0 is listed for virtual interfaces and ether stubs.

C. The number of Interrupts is listed, which indicates network efficiency.

D. A number equal to the number of transmitted Ethernet frames is listed for physical links.

E. The number of packets that were interrupted by a collision is listed, which may indicate hardware problems.

Answer: C,E

Explanation: In this output, the statistics for interrupt (INTRS) are significant. Low interrupt numbers indicate greater efficiency in performance. If the interrupt numbers are high, then you might need to add more resources to the specific link.

Example:

```
# dlstat -r -i 1
```

```
LINK IPKTS RBYTES INTRS POLLS CH<10 CH10-50 CH>50
```

```
e1000g0 101.91K 32.86M 87.56K 14.35K 3.70K 205 5
```

```
nxge1 9.61M 14.47G 5.79M 3.82M 379.98K 85.66K 1.64K
```

```
vnic1 8 336 0 0 0 0 0
```

```
e1000g0 0 0 0 0 0 0 0
```

```
nxge1 82.13K 123.69M 50.00K 32.13K 3.17K 724 24
```

```
vnic1 0 0 0 0 0 0 0
```

Note: `dlstat show-link [-r [-F] | -t] [-i interval] [-a] [-p] [-o field[,...]] [-u R|K|M|G|T|P] [link]`

Display statistics for a link.

-r

Display receive-side statistics only. Includes bytes and packets received, hardware and software drops, and so forth.

List of supported RX fields:

link

iusedby

ibytes

ipkts

intrs

polls

hdrops: hardware drops

sdrops: software drops (owing to bandwidth enforcement)

ch<10: number of packet chains of length < 10

ch10-50: number of packet chains of length between 10 and 50

ch>50: number of packet chains of length > 50

Reference: man dlstat

Question No : 27

You want to configure an iSCSI target device on your system.

Select the group package required to install this functionality on your system.

- A. storage-server
- B. solaris-small-server
- C. storage-avs
- D. storage-nas

Answer: A

Explanation: How to Create an iSCSI LUN

The disk volume provided by the server is referred to as the target. When the LUN is associated with an iSCSI target, it can be accessed by an iSCSI initiator.

The following tasks are completed on the system that is providing the storage device.

1. Install the COMSTAR storage server software.

```
target# pkg install storage-server
```

Etc.

Question No : 28

Which command should you choose to display the current parameters for the FSS scheduler?

- A. `dispadmin -c FSS`
- B. `prionctl -c FSS`
- C. `dispadmin -c FSS -g`
- D. `prionctl -c FSS -g`

Answer: C

Explanation: The `dispadmin` command displays or changes process scheduler parameters while the system is running.

-c class

Specifies the class whose parameters are to be displayed or changed. Valid class values are: RT for the real-time class, TS for the time-sharing class, IA for the inter-active class, FSS for the fair-share class, and FX for the fixed-priority class. The time-sharing and inter-active classes share the same scheduler, so changes to the scheduling parameters of one will change those of the other.

-g

Gets the parameters for the specified class and writes them to the standard output.

Reference: `man dispadmin`

Question No : 29

The following information is displayed during the migration of a ZFS file system:

DATASET	BYTES XFRD	EST BYTES LEFT	ERRORS	ELAPSED TIME
pool1/data1	409M	25.5G	-	00:07:56
pool1/data1	417M	26.0G	-	00:08:06
pool1/data1	422M	26.4G	-	00:08:16
pool1/data1	428M	26.7G	-	00:08:26
pool1/data1	433M	27.0G	-	00:08:36
pool1/data1	437M	27.3G	-	00:08:46

Which answer correctly describes the status of the migration?

- A. The new file system is named pool1/data1, the shadowstat command was used to generate the output, and the migration is still in progress.
- B. The shadowed file system is named pool1/data1, the shadowstat command was used to generate the output, and the migration is still in progress
- C. 25.5 GB of data has been transferred so far. 409 MB of data remains to be transferred
- D. The shadowed file system is named pool1/data1. the shadowstat command was used to generate the output, and the migration is still in progress with 409 MB remaining to be transferred to the target filesystem.
- E. The new file system is named pool1/data1, the shadow command was used to generate the output, and the migration is still in progress.
- F. The shadowed file system is named pool1/data1, the shadow command was used to generate the output, and the migration is still in progress.
- G. 27.3 GB is the exact size of the data that is yet to be transferred from the shadowed file system.

Answer: F

Question No : 30

Which three are true of a solaris 0 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 solaris 0 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 solaris 0 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 solarisIO branded zone.

H. The solarisIO branded zone cannot be an NFS server

Answer: B,E,H

Question No : 31

You created a new zpool. Now you need to migrate the existing ZFS file system from pool1/prod to pool2/prod.

You have these requirements:

1. Users must have access to the data during the migration, so you cannot shutdown the file system while the migration takes place.
2. Because you want to copy the data as quickly as possible, you need to increase the server resources devoted to the ZFS migration.

Which method would you use to modify the ZFS shadow migration daemon defaults to increase the concurrency and overall speed of migration?

A. Svccfg - s filesystem/shadowd:default

```
setprop config_params/shadow_threads=integer: 16
end
```

```
svcadm refresh filesystem/shadowd: default
```

B. Specify the -b <blocksize> option with the zfs create command and increase the value of <blocksize>

C. Use the -o -volblocksize=<blocksize>option with the zfs create command and increase the value of the default <blocksize>.

D. Svccfg -s filesystem/zfs: default

```
setprop config_params/shadow_threads = integer: 16
end
```

```
svcadm refresh filesystem/zfs:default
```

Answer: A

Explanation: shadowd is a daemon that provides background worker threads to migrate data for a shadow migration. A shadow migration gradually moves data from a source file system into a new “shadow” file system. Users can access and change their data within the shadow file system while migration is occurring.

The shadowd service is managed by the service management facility, smf(5). Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). The service's status can be queried using the svcs(1) command.

The svccfg(1M) command can be used to manage the following parameter related to shadowd:

config_params/shadow_threads

Note: Oracle Solaris 11: In this release, you can migrate data from an old file system to a new file system while simultaneously allowing access and modification of the new file system during the migration process.

Setting the shadow property on a new ZFS file system triggers the migration of the older data. The shadow property can be set to migrate data from the local system or a remote system with either of the following values:

file:///path

nfs://host:path

Reference: man shadowd

Question No : 32

Which three statements accurately describe the Automated Installation (AI) client?

- A.** If the AI client does not match any criteria to use a custom manifest or script, the default manifest is used.
- B.** If the AI client does not match any criteria to use a custom manifest or script, the automated installation aborts.
- C.** Any manifest or script in a service can be designated to be the default for that service.
- D.** Only the default.xml file is used as the default AT client manifest.
- E.** If a client system does not use any SC profile, then an interactive tool opens on that client at first boot after that client installation to complete the configuration of that client.
- F.** If a client system does not use any SC profile, then the install server will use the default SC profile.

Answer: A,D,E

Explanation: A: Each client uses one and only one AI manifest to complete its installation. The AI manifest is selected for a client according to the following algorithm:

* If no custom AI manifests are defined for this install service, the default AI manifest is used. The default AI manifest is not associated with any client criteria

etc.

D: When you create a new install service, `install_service_image_path/auto_install/manifest/default.xml` is the initial default AI manifest for that install service.

E: Each client can use any number of system configuration profiles. If a client system does not use any configuration profile, then an interactive tool opens on that client at first boot after that client installation to complete the configuration of that client.

Reference: Installing Oracle Solaris 11 Systems, Customizing Installations

Question No : 33

The following information is displayed about the compress/zip software package, which is currently installed on this system:

```
NAME (PUBLISHER)VERSIONINFO
```

```
Compress/zip3.1.2-0.175.0.0.0.0.537if-
```

```
NAMEVERSIONDATECOMMENT
```

```
Compress/zip3.109 Dec 2011 04:50:38 ESTNone
```

Which statement describes the information that is displayed for the compress/zip software package?

A. This package cannot be removed.

B. This package can be updated to a new version when the new version of the package becomes available.

- C. This package cannot be updated.
- D. This package can be updated to version 3.1.3 but not 3.2.
- E. This package cannot be downgraded to version 3.1.1.

Answer: D

Explanation:

An “f” in the F column indicates the package is frozen. If a package is frozen, you can only install or update to packages that match the frozen version.

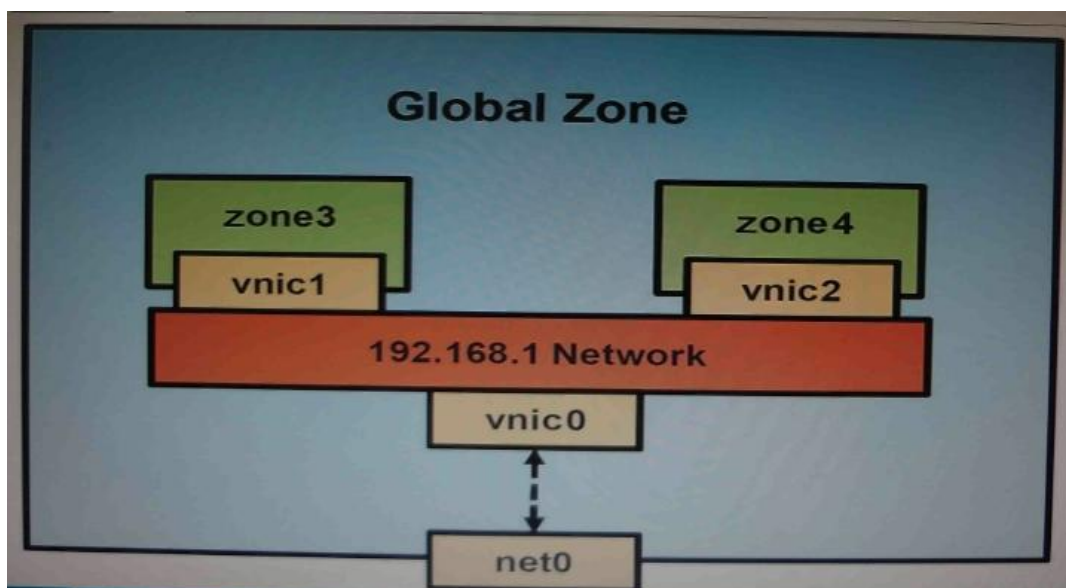
Note: The “i” in the I column indicates that these packages are installed in this image.

References: Adding and Updating Oracle Solaris 11 Software Packages, Showing Package Install State Information

Question No : 34

You have been asked to troubleshoot the initial configuration of a virtual network connecting two local zones with the outside world.

View the exhibit.



The command

```
dladm create-vnic -1 vswitch192.168.1 vnic1
```

fails with the error

```
dladm: invalid link name 'vswitch192.168.1'
```

What is the reason for this error?

- A. The name vswitch192.168.1 is not legal.
- B. The zone must be specified with
dladm create-vnic -z zone3 vnic1.
- C. The virtual interface must be specified with
dladm create-vnic -z zone3 vnic1.
- D. The virtual interface must be created with
ipadm create-vnic -1 switch192.168.1.
- E. The virtual switch must be created first with
dladm create -etherstub vswitch192.168.1.

Answer: E

Explanation: There is no data-link named vswitch192.168.

We need to create an etherstub first.

See Note and example below for details.

Note: Create a VNIC in the system's global zone.

```
# dladm create-vnic -l data-link vnic-name  
data-link is the name of the interface where the VNIC is to be configured.  
-l link, --link=link
```

link can be a physical link or an etherstub.

vnic-name is the name that you want to give the VNIC.

For example, to create a VNIC named vnic0 on interface e1000g0, you would type the following:

```
# dladm create-vnic -l e1000g0 vnic0
```

Example: Creating a Virtual Network Without a Physical NIC

First, create an etherstub with name stub1:

```
# dladm create-etherstub stub1
```

Create two VNICs with names hello0 and test1 on the etherstub. This operation implicitly creates a virtual switch connecting hello0 and test1.

```
# dladm create-vmnic -l stub1 hello0
# dladm create-vmnic -l stub1 test1
```

Reference: man dladm

Question No : 35

Review the zonestat command

```
zonestat -q -r physical-memory -R high -z dbzone -p -p "zones" 10 24h 60m
```

Select the option that correctly describes the information that is displayed by this command

- A.** It is a sample of dbzone's physical memory usage taken every hour over a 24-hour period Only the top 10 samplings of peak memory usage are displayed. All other utilization data is eliminated
- B.** It is a sample of dbzone's CPU, virtual memory, and networking utilization Physical memory is excluded from the report The sampling is taken every 10 minutes over a 24 hour period and peak utilization is displayed each hour.
- C.** It is a sample of dbzone's CPU, virtual memory, and networking utilization Physical memory is excluded from the report. The sampling is taken every 10 seconds over a 24-hour period and displayed each hour.
- D.** It is a sample of dbzone's memory and CPU utilization every 10 seconds over a 24-hour period Only peak virtual memory usage and CPU utilization are displayed each hour All other utilization data is eliminated.
- E.** It is a sample of dbzone's physical memory usage taken every 10 seconds over a 24-hour period. Only peak memory usage is displayed each hour All other utilization data is eliminated.

Answer: E

Question No : 36

you are to create a file system with the following specifications:

lzjb compression enabled

cannot consume more than 2 GB from the storage pool

redundant data at the block level eliminated

mounted as /data

Which command creates the desired file system?

- A. `zfs create -o mountpoint=/data,compression=on,algorithm=lzjb,deduplication=on,quota=2g/pool1/data`
- B. `zfs create -o mountpoint=/data compression=on algorithm=lzjb deduplication=on quota=2g pool1/data`
- C. `zfs create -o mountpoint=/data -o compression=on -o dedup=on -o quota=2g pool1/data`
- D. `zfs create -o mountpoint=/data -o compression=on -o algorithm=lzjb -o deduplication=on -o quota=2g pool 1/data`
- E. `zfs create pool 1 /data zfs set mountpoint=/data,quota=2g,dedup=on,compression=on /pool1/data`

Answer: C

Question No : 37

Which three options describe the purpose of the `zonep2vchk` command?

- A. Used on a Solaris 10 global zone to access the system for problems before migrating that system to a Solaris 10 branded zone.
- B. Used to access a Solaris 10 global zone for problems before migrating that zone to a Solaris 11 global zone
- C. Used to create `zonecfg` template for a Solaris 10 global zone that that will be migrated to a solaris10 branded zone.
- D. Used to migrate an Oracle Solaris 11 global zone to a non-global zone.
- E. Used to migrate a Solaris 10 global zone to a non-global zone on the same server; the non-global zone can then be migrated to a Solaris 11 server as a Solaris10 branded zone.

Answer: C,D,E

Explanation: `zonep2vchk`

- check a global zone's configuration for physical to virtual migration into non-global zone

The `zonep2vchk` utility is used to evaluate a global zone's configuration before the process

of physical-to-virtual (p2v) migration into a non-global zone.

The p2v process involves archiving a global zone (source), and then installing a non-global zone (target) using that archive

zonep2vchk serves two functions. First, it can be used to report issues on the source which might prevent a successful p2v migration. Second, it can output a template zonecfg, which can be used to assist in configuring the non-global zone target.

zonep2vchk can be executed on a Solaris 10 or later global zone. To execute on Solaris 10, copy the zonep2vchk utility to the Solaris 10 source global zone.

When run on Solaris 10, a target release of S11 can be specified, which will check for p2v into a Solaris 10 Branded zone.

Reference: man zonep2vchk

Question No : 38

Before booting testzone, a non-global zone, you want to connect to the zone's console so that you can watch the boot process. Choose the command used to connect to testzone's console.

- A. zoneadm -C testzone
- B. zoneadm -console testzone
- C. zlogin -z testzone console C
- D. zlogin -z testzone -C
- E. zlogin -C testzone
- F. zoneadm -z testzone -C

Answer: E

Question No : 39

If the Automated Installer is used, rather than JumpStart, to install the Oracle Solaris 11 operating system. Identify the two correctly matching pairs of equivalent functionality. (Choose two.)

A. AJ: manifest files

JumpStart: begin script

B. AI:installadmcreate-client

JumpStart: setup_install_sen/er

C. AI SMF system configuration profile files JumpStart: profile files

D. AI: SMF system configuration profile files JumpStart: finish scripts and sysidcfg files

Answer: A

Question No : 40

Review the zonestat command:

```
zonestate - q physical - memory -R high -z -p -p "zones" 10 24h 60m
```

Select the option that correctly describes the information that is displayed by this command.

A. It is a sample of dbzone's physical memory usage taken every hour over a 24-hour period.

Only the top 10 samplings of peak memory usage are displayed.

All other utilization data is eliminated.

B. It is a sample of dbzone's CPU, virtual memory, and networking utilization.

Physical memory is executed from the report.

The sampling is taken every 10 minutes over a 24-hour period and peak utilization is displayed each hour.

C. It is a sample of dbzone's CPU, virtual memory, and networking utilization.

Physical memory is executed from the report.

The sampling is taken every 10 minutes over a 24-hour period and displayed each hour.

D. It is a sample of dbzone's physical memory usage taken every 10 seconds and 24-hour period.

Only peak virtual memory usage and CPU utilization are displayed each hour.

All other Utilization data is eliminated.

E. It is a sample of dbzone's physical memory usage taken every 10 seconds and 24-hour period.

Only peak memory usage is displayed each hour.

All other utilization data is eliminated.

Answer: D

Explanation: * (Not A, B, C): interval (here 10 seconds): Specifies the length in seconds to pause between each interval report.

* duration (here 24 h)

* -R report[,report] (here high)

Print a summary report.

high Print a summary report detailing the highest usage of each resource and zone during any interval of the zonestat utility invocation.

Note: The zonestat utility reports on the cpu, memory, and resource control utilization of the currently running zones. Each zone's utilization is reported both as a percentage of system resources and the zone's configured limits.

The zonestat utility prints a series of interval reports at the specified interval. It optionally also prints one or more summary reports at a specified interval.

The default output is a summary of cpu, physical, and virtual memory utilization. The -r option can be used to choose detailed output for specific resources.

Question No : 41

The ZFS configuration on your server is:

Pool1 6.67G31K/pool

Pool1/data31K31K/data

Select the three commands that you would use to 1. Create, 2. List, and 3. Delete a snapshot of the /data file system.

- A. zfs snapshot pool1/data@now
- B. zfs create snapshot pool1/data@now
- C. zfs list -t snapshot
- D. zfs list -t snapshot pool1/data
- E. zfs destroy pool1/data@now
- F. zfs destroy snapshot pool1/data@now

Answer: A,D,E

Explanation: A: Snapshots are created by using the zfs snapshot command, which takes