# Oracle OPN 1Z0-820 Practice Questions

**Shared by Noble on 10-11-2026**

**For More Free Questions and Preparation Resources**

Check the Links on Last Page

## Question 1

Question Type: MultipleChoice

You are troubleshooting the Oracle Solaris11 Automated Installer (AI), which is not connecting with the IPS software repository.

Which three steps will help determine the cause of DNS name resolution failure?

### Options:

A- Verify the contents of /etc/resolve.conf.
B- Run netstat -nr to verify the routing to the DNS server.
C- Ping the IP address of the IPS server to verify connectivity.
D- On the installation server, verify that the menu.1st file for the client points to a valid boot arc hive.
E- Run df -k to verify that the boot directory containing the boot archive is loopback mounted under /etc/netboot.
F- Run the command /sbin/dhcpinfo DNSserv to ensure that the DHCP server providing the DNS server information.

### Answer:

A, B, F

### Explanation:

Check DNS

* (A) Check whether DNS is configured on your client by verifying that a non-empty /etc/resolv.conf file exists.

* (F) If /etc/resolv.conf does not exist or is empty, check that your DHCP server is providing DNS server information to the client:

# /sbin/dhcpinfo DNSserv

If this command returns nothing, the DHCP server is not set up to provide DNS server information to the client. Contact your DHCP administrator to correct this problem.

* (B) If an /etc/resolv.conf file exists and is properly configured, check for the following possible problems and contact your system administrator for resolution:

** The DNS server might not be resolving your IPS repository server name.

** No default route to reach the DNS server exists.

## Question 2

Question Type: MultipleChoice

You create a flash archive of the Solaris 10 global zone on the serves named sys

### Options:

A- zoneadm -z s10 -zone install - a /net/backup_server/s10-system.flar -u
A- The archive name is s10-system.flar, and it is stored on a remote server named backup_server.
On sysA, you create a Solaris 10 branded zone named s10-zone.
You want to use the flash archive, located On' /net/bactup_servers/10-system.flar, to install the Operating system in the s10-zone zone.
Which command do you choose to install the s10-system.flar archive in the Solaris 10 branded zone (s10-zone)?
B- zonecfg -z s10 -zone install - a /net/backup_server/s10-system.flar -u
C- zoneadm - z s10 -zone clone - s /net/backup_server/s10-system.flar
D- zone cfg - a s10-zone create - t SUNWsolaris10\
</net/backup_server/s10-system.flar
E- zonecfg -z s10-zone install -f /net/backup/backup_server/s10-system.flar

### Answer:

A, A

### Explanation:

The zoneadm command is the primary tool used to install and administer non-global zones. Operations using the zoneadm command must be run from the global zone on the target system.

How to Install the solaris10 Branded Zone

A configured solaris10 branded zone is installed by using the zoneadm command with the install subcommand.

Example: global# zoneadm -z s10-zone install -a /net/machine_name/s10-system.flar --u

## Question 3

Question Type: MultipleChoice

You are configuring NFS on a server. Choose the two statements that are true.

## Options:

A- Resources listed in /etc/dfs/dfstab are automatically shared on boot up.
B- A directory cannot be shared if a subdirectory below it is already shared.
C- Renaming a share created with the zfs set share command is not supported.
D- NFS and SMB protocols cannot be used simultaneously to share the same directory.

## Answer:

A, C

## Explanation:

A: ZFS can automatically share file systems by setting the sharenfs property. Using this property, you do not have to modify the /etc/dfs/dfstab file when a new file system is shared. The sharenfs property is a comma-separated list of options to pass to the share command. The value on is an alias for the default share options, which provides read/write permissions to anyone. The value off indicates that the file system is not managed by ZFS and can be shared through traditional means, such as the /etc/dfs/dfstab file. All file systems whose sharenfs property is not off are shared during boot.

Incorrect answers:

D: A share can define options for both NFS and SMB sharing.

# Question 4

Question Type: MultipleChoice

A non-global zone named testzone is currently running.

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

## Options:

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 5

Question Type: MultipleChoice

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

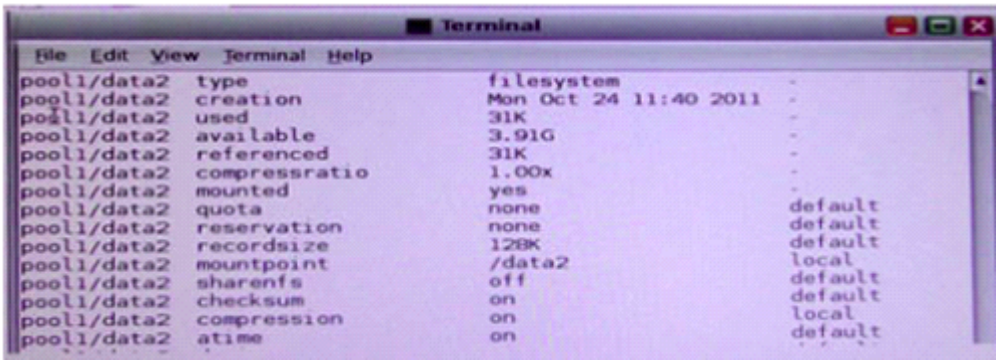


The application development team requested an up-to-date copy of the data from the /prod_data file system. You decide to give the team one of the disk drives containing the data by breaking the mirror, removing the disk c4t1d0 and mounting c4t1 dO under a new mount point named /dev_data. Identify the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

## Options:

A- zfs destroy pooh /prod_data
zfs create pool1/prod_data c4t1d0 zfs create pool1/dev_data c4t1d0
B- zfssplitpool1/prod_data
zfs mount -F zfs pool1/dev_data c4t1d0
C- zpool split pooll 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
3 e zfs split pool1/prod_data -n pool2/dev_data zfs set mountpoint=/dev_data pool2/dev_data

### Answer:

C

## Question 6

Question Type: MultipleChoice

Your task is to convert a JumpStart sysidcfg file to an Automated Installer (Al) sc_profile.xml file, using js2ai. Select two unsupported items that will require changes. (Choose two.)

### Options:

A- terminal=xterms
B- name_service=NIS+
C- timezone=US/Pacific
D- system_locale=en_US
E- network_interface=PRIMARY
F- root_password=rJmv5LUXM10cU

### Answer:

A, D

## Question 7

Question Type: MultipleChoice

In an effort to reduce storage space on your server, you would like to eliminate duplicate copies of data in your server's ZFS file systems.

How do you specify that pool1/data should not contain duplicate data blocks (redundant data) on write operations?

### Options:

A- zfs create - o compression=on pool1/data
B- zpool create -o deduplication =on pool1; zfs create pool1/data
C- zfs create - o deduplication=on pool1; zfs create pool1/data
D- zfs create - o dedupratio=2 pool1/data
E- zfs create - o dedup=on pool1/data

## Answer:

E

## Explanation:

ZFS Deduplication Property

Solaris Express Community Edition, build 129: In this Solaris release, you can use the deduplication property to remove redundant data from your ZFS file systems. If a file system has the dedup property enabled, duplicate data blocks are removed synchronously. The result is that only unique data is stored and common components are shared between files.

You can enable this property as follows:

# zfs set dedup=on tank/home

# Question 8

Question Type: MultipleChoice

Which three are true of a solansl 0 branded zone running on an Oracle Solaris 11 system? (Choose three.)

## Options:

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 solarisl 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 solarisl 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 solarislO branded zone.
H- The solarisl0 branded zone cannot be an NFS server

## Answer:

B, E, H

## Question 9

Question Type: MultipleChoice

You need to set up a local package repository to serve 75 client systems. Multiple clients will being the package repository concurrently and you need to ensure that the local repository performs very well under this heavy load, especially during package intensive operations.

Which option would ensure the best performance of the repository during package-intensive rations by multiple clients?

### Options:

A- Set up multipathing on the package repository server to distribute the network load multiple network interfaces.
B- Deploy a second instance of the package repository server to run as a read writable mirror.
C- Deploy a second instance of the package repository server to run as a read-only mirror.
D- Deploy a second instance of the package repository server to run as a clone of the primary repository server.
E- Deploy a package repository locally on each client.

### Answer:

A

## Question 10

Question Type: MultipleChoice

You need to set up a local package repository to serve 75 client systems. Multiple clients will be accessing the package repository concurrently and you need to ensure that the local repository performs very well under this heavy load, especially during package-intensive operations. Which option would ensure the best performance of the repository during package-intensive operations by multiple clients?

### Options:

A- Set up multipathing on the package repository server to distribute the network load across multiple network interfaces.
B- Deploy a second instance of the package repository server to run as a read-writable mirror.
C- Deploy a second instance of the package repository server to run as a read-only mirror.
D- Deploy a second instance of the package repository server to run as a clone of the primary repository server

E- Deploy a package repository locally on each client.

## Answer:

C

## Question 11

Question Type: MultipleChoice

How should you permanently restrict the non-global zone testzone so that it does not use more than 20 CPU shares while it is running?

## Options:

A- While configuring the zone, add this entry:
add rct1
set name = capped.cpu-shares
add value (priv = privileged, limit = 20, action = none)
end
exit
B- While configuring the zone, add this entry:
add rct1
set name= zone.cpu-shares
add value (priv=privileged, limit=20, action=none)
end
exit
from command line, enter:
# dispadmin - d FSS
C- From the command line enter:
#prct1 -n zone.cpu-shares - r - v 20 - i zone testzone
D- From the command line, enter:
#prct1 - n zone.cpu-shares - v 80 - r - i zone global

## Answer:

C

## Explanation:

Theprctlutility allows the examination and modification of the resource controls associated with an active process, task, or project on the system. It allows access to the basic and privileged limits and the current usage on the specified entity.

How to Change thezone.cpu-sharesValue in a Zone Dynamically

This procedure can be used in the global zone or in a non-global zone.

Be superuser, or have equivalent authorizations.

For more information about roles, seeConfiguring and Using RBAC (Task Map) inSystem Administration Guide: Security Services.

Use theprctlcommand to specify a new value forcpu-shares.

# prctl -n zone.cpu-shares -r -v value -i zone zonename

idtypeis either thezonenameor thezoneid.valueis the new value.

Note: project.cpu-shares

Number of CPU shares granted to a project for use with the fair share scheduler

## Question 12

Question Type: MultipleChoice

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.

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

### Options:

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/vfstabfile (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.

Incorrect answers:

A: not a network connectivity problem

D: The sharemgr command configures share groups and the shares contained within them.

## Question 13

Question Type: MultipleChoice

You upgraded your server to Oracle Solaris 11 and you imported zpool (pool1) that was created in Solaris 10. You need to create an encrypted ZFS file system in pool1, but first you need to make sure that your server supports ZFS encryption.

Which four statements are true for support of ZFS encryption?

## Options:

A- The encrypted file system must have been created in Oracle Solaris11. To encrypt a ZFS file system from a previous version of Solaris, upgrade the zpool and create a new encrypted ZFS file system into the encrypted ZFS file system.
B- If you plan to create an encrypted file system in an existing zpool, the zpool must be upgraded to ZFS version 30.
C- ZFS encryption is integrated with the ZFS command set and no additional packages need to be installed.
D- ZFS encryption requires that the ZFS Dataset Encryption package be installed.
E- If you plan to create an encrypted file system in an existing zpool, the pool must be upgraded to ZFS version 21, minimum.
F- Encryption is supported at the pool or dataset (file system) level.
G- Encryption is supported at the pool level only for every file system in the pool will be encrypted.
H- You cannot create an encrypted file system in a zpool that was created prior to oracle Solaris11. Create a new zpool in Solaris11, create an encrypted ZFS file system in the new zpool, and move or copy the data from the existing file system into the new encrypted file system.

## Answer:

A, B, C, F

## Explanation:

A (not H): You can use your existing storage pools as long as they are upgraded. You have the flexibility of encrypting specific file systems.

B(not E): Can I enable encryption on an existing pool?

Yes, the pool must be upgraded to pool version 30 to allow encrypted ZFS file systems and volumes.

C(not D): ZFS encryption is integrated with the ZFS command set. Like other ZFS operations, encryption operations such as key changes and rekey are performed online.

F (not G): Encryption is the process in which data is encoded for privacy and a key is needed by the data owner to access the encoded data. You can set an encryption policy when a ZFS dataset is created, but the policy cannot be changed.

To Get Premium Files for 1Z0-820 Visit

https://www.p2pexams.com/products/1z0-820

For More Free Questions Visit

https://www.p2pexams.com/oracle/pdf/1z0-820

