

Oracle

Exam 1z0-820

Upgrade to Oracle Solaris 11 System Administrator

Version: 7.0

[Total Questions: 133]

Question No : 1

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 : 2

For an Oracle Solaris 11 Automated Installation (AI), select the two properties that can be configured using System Configuration profiles (SC profiles). (Choose two.)

- A. IP address of the AI server
- B. passwords for user accounts
- C. target disk slice for installation
- D. NWAM active network configuration profile (NCP)
- E. IP address of the IPS (Image Packaging System) repository
- F. SI_MANIFEST_SCRIPT, which defines the URL of the manifest script

Answer: B,F

Question No : 3

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 : 4

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 : 5

The resource control in your zone is:

```
rctl
```

```
name: zone.cpu-shares
```

```
value: (pvn=phvJleged,limit=20,action=none)
```

When you boot the zone, this message is displayed:

```
WARNING: The zone.cpu-shares rctl is set but FSS is not the default scheduling class for this zone. FSS will be used for processes in the zone but to get the full benefit of FSS, it should be the default scheduling class.
```

Which option will resolve this issue?

- A. in the global zone, run this command to change the process scheduler for the global zone `dispadm -d FSS`
- B. in the non-global zone, run this command to change the process scheduler for this specific zone: `dispadm -d FSS`
- C. Change the zone resource control to:
value: (priv=privileged,limit=20,action=none,default)
- D. Use `svccfg` to modify the `general/scheduler` property in the `svc:/system/zones:default` service:
`svccfg -s system/zones:default editprop general/scheduler=FSS`

Answer: D

Question No : 6

You are setting up an automated installer (AI) install server and issue the following command:

```
installadm create-service -n prod_ai -s /repo/prod_ai.iso \  
-i 192.168.1.100 -c 5 -d /export/repo
```

Which four options describe the install server that you have configured?

- A. The service name is `prod_ai`.
- B. DHCP base IP address is 192.168.1.100
- C. The initial IP address for the install clients will be 192.168.1.100. This IP address is temporary. After the client is booted, it will use IP addresses in the following range: 192.168.1.101-105.
- D. Five IP addresses are allocated for DHCP clients, starting with 192.168.1.100.
- E. The Install server will support up to five clients.
- F. The AI net image ISO file is located in `/repo/prod` and the net image ISO will be unpacked in `/export/repo`.
- G. The AI net image ISO file is located in `/repo/repo` and is named `/repo/prod/_ai.iso`.

Answer: A,B,D,F

Explanation: A: `-n <svcname>`

Uses this install service name instead of default service name.

B: -i <dhcp_ip_start>

Sets up a new DHCP server. The IP addresses, starting from dhcp_address_start, are set up.

D: -c <count_of_ipaddr>

Sets up a total number of IP addresses in the DHCP table equal to the value of the count_of_ipaddr. The first IP address is the value of dhcp_ip_start that is provided by the -i option.

F: -s <srcimage>

Specifies location of AI ISO image to use for setting up the install service.

<targetdir>

Required: Specifies location to set up net image.

Reference: man installadm

Question No : 7

You have already generated a 256-bit AES raw key and named the keystore file /mykey. You need to use the key to create an encrypted file system.

Which command should you use to create a ZFS encrypted file system named pool1/encrypt using the /mykey keystore?

- A. zfs create - o encryption = /mykey pool1/encrypt
- B. zfs create - o encryption = 256-ccm - o keysource = raw, file : ///my key pool1/encrypt
- C. zfs create - o encryption = AES keysource = /mykey pool1/encrypt
- D. zfs create - o encryption = on keystore = /mykey pool1/encrypt

Answer: B

Explanation: Example: Encrypting a ZFS File System by Using a Raw Key

In the following example, an aes-256-ccm encryption key is generated by using the pktool command and is written to a file, /cindykey.file.

```
# pktool genkey keystore=file outkey=/cindykey.file keytype=aes keylen=256
```

Then, the /cindykey.file is specified when the tank/home/cindy file system is created.

```
# zfs create -o encryption=aes-256-ccm -o keysource=raw,file:///cindykey.file  
tank/home/cindys
```

Reference: Oracle Solaris ZFS Administration Guide, Examples of Encrypting ZFS File Systems

Question No : 8

Which three statements are true for the Oracle Solans 111image Packaging System (IPS)?
(Choose three.)

- A. a local repository should be stored on a 2FS file system.
- B. Network access to an IPS repository is provided to client systems using NFS.
- C. Network access to an IPS repository is provided to client systems using NFS and HTTP.
- D. The distribution constructor can get packages from a local IPS to build an installation image.
- E. A mirror repository contains all of the metadata and content for software packages.
- F. Clients access a mirror repository to obtain a publisher's catalog and software packages.

Answer: D,E,F

Question No : 9

New features were added to ZFS in Oracle Solaris11. Your justification to upgrade from Solaris10 to oracle Solaris11 is that it will be possible to take advantage of the enhancements that were made to ZFS.

Identify the three ZFS functions and features that are included in Oracle Solaris 11, but not in Solaris 10.

- A. Encrypted ZFS datasets

- B. Ability for ZFS to detect and remove redundant data from the tile system
- C. Shadow Data Migration
- D. Ability to split a mirrored ZFS storage pool
- E. Ability to use ZFS on the boot drive and boot to a ZFS root file system.
- F. elimination of the swap file system when using ZFS on the root disk

Answer: A,B,E

Explanation: A: ZFS encryption was introduced in ZFS Pool Version Number 30, Solaris Nevada b149.

Filesystem encryption since Solaris 11 Express

B: Deduplication was introduced in ZFS Pool Version Number 21, Solaris Nevada b128.

E: Booting From a ZFS Root File System

Both SPARC based and x86 based systems use the new style of booting with a boot archive, which is a file system image that contains the files required for booting. When a system is booted from a ZFS root file system, the path names of both the boot archive and the kernel file are resolved in the root file system that is selected for booting.

Reference: Oracle Solaris ZFS Administration Guide, Booting From a ZFS Root File System

Question No : 10

Which three Installation option allow for a "hands free" and "unattended" Installation of the Solaris 11 environment?

- A. Jumpstart
- B. LiveCD
- C. A text Installation over the network
- D. An Automated Installation performed on an x86 client
- E. An Automated Installation using media from a local DVD or USB drive
- F. An Automated Installation using a networked repository

Answer: D,E,F

Explanation: Oracle Solaris 11 uses Automated Installer (AI) for unattended installations.

Unattended installations are possible by placing the contents of the AI Image media (or ISO image contents from a download) on an AI server.

Reference: Differences between Oracle Solaris 10 and 11 for System Administrators

<http://www.oracle.com/technetwork/server-storage/solaris11/overview/solaris-matrix-1549264.html>

Question No : 11

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 : 12

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