

100% Money Back
Guarantee

Vendor:Oracle

Exam Code:1Z0-821

Exam Name:Oracle Solaris 11 System Administration

Version:Demo

QUESTION 1

Which network protocol provides connectionless, packet-oriented communication between applications?

- A. TCP
- B. UDP
- C. IP
- D. ICMP
- E. NFS
- F. IPSec

Correct Answer: B

The User Datagram Protocol (UDP) is one of the core members of the Internet Protocol Suite, the set of network protocols used for the Internet. With UDP, computer applications can send messages, in this case referred to as datagrams, to other hosts on an Internet Protocol (IP) network without requiring prior communications to set up special transmission channels or data paths. Compared to TCP, UDP is a simpler message-based connectionless protocol.

QUESTION 2

The su command by default makes an entry into the log file for every su command attempt.

The following is a single line from the file:

```
SU 12/18 23:20 + pts/1 user1-root
```

What does the + sign represent?

- A. unsuccessful attempt
- B. successful attempt
- C. The attempt was from a pseudo terminal, and not the console.
- D. The attempt was from a user that is in the adm group, same as root.
- E. Time zone is not set.

Correct Answer: B

The sulog file, `/var/adm/sulog`, is a log containing all attempts (whether successful or not) of the su command. An entry is added to the sulog file every time the su command is executed. The fields in sulog are: date, time, successful (+) or unsuccessful (-), port, user executing the su command, and user being switched to. In the preceding example, all su attempts were successful, except for the attempt on 2/23 at 20:51, when user pete unsuccessfully attempted to su to user root.

Look for entries where an unauthorized user has used the command inappropriately. The following entry shows a successful (indicated by +) su from user userid to root.

QUESTION 3

You are going to create live zones on your server. Disk space is critical on this server so you need to reduce the amount of disk space required for these zones. Much of the data required for each of these zones is identical, so you want to eliminate the duplicate copies of data and store only data that is unique to each zone.

Which two options provide a solution for eliminating the duplicate copies of data that is common between all of these zones?

- A. Create the zones by using sparse root zones.
- B. Set the dedup property to on and the dedupratio to at least 1.5 for the zpool. Create a separate ZFS file system for each zone in the zpool.
- C. Put all of the zones in the same ZFS file system and set the dedupratio property for the ZFS file system to at least 1.5.
- D. Put all of the zones in the same ZFS file system and set the dedup property for the file system to on.
- E. Put each zone in a separate ZFS file system within the same zpool. Set the dedup property to on for each ZFS file system.

Correct Answer: DE

In Oracle Solaris 11, you can use the deduplication (dedup) 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.

QUESTION 4

The current ZFS configuration on your server is:

```
pool1 124K 3.91G 32K /pool1
```

```
pool1/data 31K 3.91G 31K /data
```

You need to create a new file system named /data2. /data2 will be a copy of the /data file system.

You need to conserve disk space on this server whenever possible.

Which option should you choose to create /data2, which will be a read/writeable copy of the /data file system, while minimizing the amount of total disk space used in pool1?

- A. `zfs set mountpoint=/data2 compression=on pool1/data2`
- B. `zfs snapshot pool1/data@now`
`zfs set mountpoint=/data2, compression=on pool1/data@now`
- C. `zfs create snapshot pool1/data@now`
`zfs send pool1/data@now | zfs recv pool1/data2`
- D. `zfs create snapshot pool1/data@now`
`zfs clone -o mountpoint=/data2 pool1/data@now pool1/data2`

E. `zfs snapshot pool1/data@nowzfs clone -o mountpoint=/data2 -ocompression=on pool1/data@now pool1/data2`

F. `zfs snapshot pool1/data@nowzfs clone -o mountpoint=/data2 pool1/data@now pool1/data2`

Correct Answer: E

`zfs snapshot [-r] [-o property=value] ... filesystem@snapname|volume@snapname` Creates a snapshot with the given name. All previous modifications by successful system calls to the file system are part of the snapshot

`zfs clone [-p] [-o property=value] ... snapshot filesystem|volume` Creates a clone of the given snapshot.

Note:

Because snapshots are fast and low overhead, they can be used extensively without great concern for system performance or disk use .

With ZFS you can not only create snapshot but create a clone of a snapshot. A clone is a writable volume or file system whose initial contents are the same as the dataset from which it was created. As with snapshots, creating a clone is nearly instantaneous, and initially consumes no additional disk space. In addition, you can snapshot a clone.

A clone is a writable volume or file system whose initial contents are the same as the original dataset. As with snapshots, creating a clone is nearly instantaneous, and initially consumes no additional space.

Clones can only be created from a snapshot. When a snapshot is cloned, it creates an implicit dependency between the parent and child.

QUESTION 5

Review the boot environments displayed on your system:

BE	Active	Mountpoint	Space	Policy	Created
oldBE	-	-	149.0K	static	2011-11-28 15:15
newBE	-	-	363.05M	static	2011-11-28 14:47
solaris	-	-	100.68M	static	2011-11-20 18:09
solaris-1	NR	/	19.07G	static	2012-01-22 07:23

Which option describes the solaris-1 BE?

- A. It is active on the next reboot.
- B. It is active now.
- C. It is inactive.
- D. It is unbootable.
- E. It is active now and on reboot.
- F. It has been removed and will no longer be available after the next reboot.

Correct Answer: E

In the below output, NR (now running) means the BE is active now and will be the active BE on reboot.

Example:

Display your existing BE information.

```
# beadm list
```

```
BE Active Mountpoint Space Policy Created
```

```
solaris NR / 12.24G static 2011-10-04 09:42
```

QUESTION 6

Which two statements are true concerning the creation of user accounts by using the useradd command?

- A. By default, it will create the user's home directory.
- B. New user accounts are unlocked but must change their password at their first login.
- C. New user accounts are in a pending activation state until a password is assigned to them.
- D. By default, a new group will be added for each new user account.
- E. By default, the UID of a new user account will be the next available number above the highest number currently assigned.
- F. By default, the UID of a new user account will be the lowest available unused number for nonsystem accounts.

Correct Answer: CE

QUESTION 7

You have connected a new printer at a fixed IP address.

It appears to work correctly most of the time, but at other times does not respond.

You suspect that the assigned address may not be unique within the network.

What command will be useful to confirm this?

- A. arp
- B. netstat
- C. ipadm show-if
- D. dladm show-addr
- E. ipadm show-addr

Correct Answer: E

`ipadm show-addr` displays all the configured addresses on the system.

Example: # ipadm show-addr ADDR OBJ TYPE STATE ADDR lo0/v4 static ok 127.0.0.1/8 lo0/v6 static ok ::1/128

QUESTION 8

Which two options are characteristics of a fast reboot?

- A. A fast reboot bypasses grub.
- B. A fast reboot cannot be used after a system panic on the x86 platform.
- C. A fast reboot can only be executed on the SPARC platform when the `config/fastreboot_default` property for the `svc:/system/boot-config:default` service is set to true.
- D. A fast reboot uses an in-kernel boot loader to load the kernel into memory.
- E. A fast reboot is the default on all platforms.

Correct Answer: CD

C: To change the default behavior of the Fast Reboot feature on the SPARC platform, so that a fast reboot is automatically performed when the system reboots, see below.

The following example shows how to set the property's value to true on the SPARC platform, so that a fast reboot is initiated by default:

```
# svccfg -s "system/boot-config:default" setprop config/fastreboot_default=true # svcadm refresh svc:/ system/boot-config:default
```

D: Fast Reboot implements an in-kernel boot loader that loads the kernel into memory and then switches to that kernel. The firmware and boot loader processes are bypassed, which enables the system to reboot within seconds. The Fast Reboot feature is managed by SMF and implemented through a boot configuration service, `svc:/ system/boot-config`. The boot-config service provides a means for setting or changing the default boot configuration parameters. When the `config/fastreboot_default` property is set to true, the system performs a fast reboot automatically, without the need to use the `reboot -f` command. This property's value is set to true on the x86 platform. For task-related information, including how to change the default behavior of Fast Reboot on the SPARC platform, see *Accelerating the Reboot Process on an x86 Based System*.

Note: One new feature, called Fast Reboot, will allow the system to boot up without doing the routine set of hardware checks, a move that can make system boot times up to two- and-a-half times faster, Oracle claimed. This feature can be handy in that an administrator applying a patch or software update across thousands of Solaris deployments can reboot them all the more quickly.

QUESTION 9

You have a process called `bigscript`, and you need to know the PID number for this process.

Which command will provide that information?

- A. `pkill bigscript`

- B. ps bigscript
- C. pgrep bigscript
- D. prstat bigscript

Correct Answer: C

Pgrep takes a process name and return a PID.

Note: pgrep looks through the currently running processes and lists the process IDs which matches the selection criteria to stdout. All the criteria have to match. For example, pgrep -u root sshd will only list the processes called sshd AND owned by root.

Incorrec answers:

ps bigscript: You can't pass a name to ps, it interprets it as arguments.

QUESTION 10

Your server has a ZFS storage pool that is configured as follows:

```
state: ONLINE
scan: none requested
config:
```

NAME	STATE	READ	WRITE	CKSUM
pool1	ONLINE	0	0	0
c3t3d0	ONLINE	0	0	0
c3t4d0	ONLINE	0	0	0

The server has two spate 140-GB disk drives: c3t5d0 c3t6d0 Which command would add redundancy to the pool1 storage pool?

- A. zpool attach pool1 c3t5d0 c3t6d0
- B. zpool attach pool1 c3t3d0 c3c5d0; zpoo1 attach pool1 c3t4d0 c3t6d0
- C. zpool mirror pool1 c3t5d0 c3t6d0
- D. zpool add pool1 mirror c3t5d0 c3t6d0
- E. zpool add raidz pool1 c3t5d0 c3t6d0

Correct Answer: A

You can convert a non-redundant storage pool into a redundant storage pool by using the zpool attach command. 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.

QUESTION 11

Which best describes the `svc:/system/boot-config` service?

- A. It is used to change the milestone on a system.
- B. It is used to set the default run level of the system.
- C. It provides the parameters used to set the system to automatically perform a fast or slow reboot.
- D. When the service is enabled, the system performs a fast reboot by default; when it is disabled the system performs a slow reboot by default.

Correct Answer: C

Starting with the Oracle Solaris 11 Express release, Fast Reboot is supported on the SPARC platform, as well as the x86 platform. On both platforms, this feature is controlled by the SMF and implemented through a boot configuration service, `svc:/system/boot-config`. The boot-config service provides a means for setting or changing the default boot configuration parameters.

The `fastreboot_default` property of the boot-config service enables an automatic fast reboot of the system when either the `reboot` or the `init 6` command is used. When the `config/fastreboot_default` property is set to `true` the system automatically performs a fast reboot, without the need to use the `reboot -f` command. By default, this property's value is set to `false` on the SPARC platform and to `true` on the x86 platform.

QUESTION 12

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 deviceset match=/dev/rdisk/ c4t0d0` 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 non-global zone, make the following entry: `add datasetset name=pool2` `add fsset dir=pool1` `set special=pool1` `set type=zfs` `pool1` 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 datasetset name=pool2` 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 fsset dir=/dataset` `special=pool2/dataset` `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.

Correct Answer: C

<http://docs.oracle.com/cd/E19253-01/819-5461/gbbst/index.html>