

# Oracle

## Exam Questions 1Z0-821

Oracle Solaris 11 System Administrator



### NEW QUESTION 1

View the exhibit to inspect the file system configuration on your server.

NAME	USED	AVAIL	REFER	MOUNTPOINT
pool1	134K	3.91G	32K	/pool1
pool1/data	31K	3.91G	31K	/data
remote	124K	3.91G	32K	/remote
remote/backup	31K	3.91G	31K	/remote/backup
rpool	11.6G	4.02G	34.5K	/rpool
rpool/ROOT	9.95G	4.02G	31K	legacy
rpool/ROOT/solaris	9.95G	4.02G	9.71G	/
rpool/dump	630M	4.04G	611M	-
rpool/export	6.07M	4.02G	32K	/export
rpool/export/home	6.04M	4.02G	32K	/export/home

View the Exhibit to inspect the file system configuration on your server.

Your department's backup policy is to perform a full backup to a remote system disk on Saturday.

On Sunday through Friday, you are to perform a differential backup to the same remote system disk:

Following your company policy, which option describes a valid procedure for backing up the /data file system to a remote disk named /remote/backup?

- A) On Saturday:  

```
zfs snapshot pool1/data@sat
zfs send pool1/data@sat > /remote/backup/full
```

 On each weekday:  
 Remove the previous daily snapshot.  

```
zfs snapshot pool1/data@daily
zfs send -i pool1/data@sat pool1/data@daily > /remote/backup/full
```
- B) On Saturday:  

```
zfs create snapshot pool1/data@sat
zfs send pool1/data@sat |zfs recv remote/backup/`date +%m%d%y`
```

 On each weekday:  
 Remove the previous daily snapshot.  

```
zfs create pool1/data@daily
zfs send -i pool1/data@sat pool1/data@daily |zfs recv remote/backup/`date +%m%d%y`
```
- C) On Saturday:  

```
zfs snapshot pool1/data@sat
zfs send pool1/data@sat > /remote/backup/full
```

 On each weekday:  
 Remove the previous daily snapshot.  

```
zfs snapshot pool1/data@daily
zfs send -i pool1/data@sat pool1/data@daily > /remote/backup/`date +%m%d%y`
```
- D) On Saturday:  

```
zfs create snapshot pool1/data@sat
zfs send pool1/data@sat | zfs recv remote/backup
```

 On each weekday:  
 Remove the previous daily snapshot.  

```
zfs create -i pool1/data@sat pool1/data@daily
zfs send pool1/data@daily |zfs recv remote/backup
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: C

### NEW QUESTION 2

dbzone is currently running on your server.

Which two methods would you use to safely and cleanly shut down dbzone and all of its applications?

- A. zlogin -z dbzone halt
- B. zoneadm -z dbzone shutdown -i0
- C. zoneadm -z dbzone shutdown
- D. zoneadm -z dbzone halt
- E. zlogin dbzone shutdown -i0

Answer: DE

**Explanation:** D: zoneadm halt command halts the specified zones. halt bypasses running the shutdown scripts inside the zone. It also removes run time resources of the zone.

E: Use: zlogin zone shutdown

to cleanly shutdown the zone by running the shutdown scripts.

Use this procedure to cleanly shut down a zone.

1. Become superuser, or assume the Primary Administrator role.

2. Log in to the zone to be shut down, for example, my-zone, and specify shutdown as the name of the utility and init 0 as the state global# zlogin my-zone shutdown -y -g0 -i 0

### NEW QUESTION 3

You attempted to reboot a system via the init command, however the system did not perform boot sequence into the Oracle Solaris Operating Environment. You are presented with a prompt from the OpenBoot PROM. Which command would you enter, to boot the system from the default device?

- A. boot -net install
- B. boot
- C. boot -default
- D. boot -s0

**Answer: B**

#### **Explanation: Boot**

With this form, boot loads and executes the program specified by the default boot arguments from the default boot device

Note: boot has the following general format: boot [device-specifier] [arguments]

where device-specifier and arguments are optional.

### NEW QUESTION 4

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).

### NEW QUESTION 5

Your users are experiencing delay issues while using their main application that requires connections to remote hosts. You run the command uptime and get the following output:

1:07am up 346 day(s), 12:03, 4 users, load average: 0.02, 0.02, 0.01 Which command will be useful in your next step of troubleshooting?

- A. ipadm
- B. traceroute
- C. dladm
- D. snoop
- E. arp

**Answer: B**

#### **Explanation: Test the remote connection with traceroute.**

The Internet is a large and complex aggregation of network hardware, connected together by gateways. Tracking the route one's packets follow (or finding the miscreant gateway that's discarding your packets) can be difficult. traceroute utilizes the IP protocol 'time to live' field and attempts to elicit an ICMP TIME\_EXCEEDED response from each gateway along the path to some host.

This program attempts to trace the route an IP packet would follow to some internet host by launching UDP probe packets with a small ttl (time to live) then listening for an ICMP "time exceeded" reply from a gateway.

### NEW QUESTION 6

Your task is to convert a JumpStart sysidcfg file to an Automated Installer (AI) sc\_profile.xml file, using js2ai.

Select two unsupported items that will require changes.

- A. terminal = zterms
- B. name\_service=NTS+
- C. timezone=US/pacific
- D. system\_locale=en\_US
- E. network\_interface=PRIMARY
- F. root\_password=rJmvLUXM10cU

**Answer: AD**

#### **Explanation: A: terminal**

The js2ai tool does not perform any translation. Make sure the terminal type specified in the sysidcfg file is supported in Oracle Solaris 11.

D: system\_locale

The js2ai tool does not perform any translation. Make sure the locale specified in the sysidcfg file is supported in Oracle Solaris 11.

### NEW QUESTION 7

user1 has a disk quota of 0.5 MB. The user attempts to run the following command on a file called .bigfile that is 495 KB in size:

```
cp bigfile /tmp
```

Will the command execute successfully?

- A. Ye
- B. Quotas do not include any of the system files such as /tmp /swap.
- C. Ye
- D. The quota is set at the directory level, not the user level.
- E. N
- F. The command will fail because it will cause him to exceed his user quota.
- G. N
- H. A user cannot place files into the /tmp directory.

**Answer: A**

**Explanation:** UFS quotas enable system administrators to control the size of file systems. Quotas limit the amount of disk space and the number of inodes, which roughly corresponds to the number of files, that individual users can acquire. For this reason, quotas are especially useful on the file systems where user home directories reside. As a rule, the public and /tmp file systems usually do not benefit significantly by establishing quotas. Note: The cp command copies files and directories.

### NEW QUESTION 8

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: D**

**Explanation:** ai\_manifest

- Automated installation manifest file format

Synopsis

/usr/share/install/ai.dtd.1

Some customizations have been made, such as the selection of specific locales.

### NEW QUESTION 9

You are troubleshooting interface net3 and you enter the following sequence of commands:

```
Command:
dladm show-if | grep net3
Output:
net3 Ethernet up 1000 full
```

```
Command:
ipadm show-if
Output:
IFNAME CLASS STATE ACTIVE OVER
net3 ip down no --
```

```
Command:
ipadm up-addr net3/v4
Output:
ipadm: cannot mark the address up: Object not found
```

Your next command should be:

- A. ipadm up-addr net3/v4
- B. ipadm enable-if -T net3
- C. <ipadm create-vnic -a 192.168.1.25/24 net3/v4
- D. ipadm create-ip -T static -a 192.168.1.25/24 -n net3
- E. ipadm create-addr -T static -a 192.168.1.25/24 net3

**Answer: E**

**Explanation:** If you are assigning a static IP address, use the following syntax:

```
# ipadm create-addr -T static -a address addrobj
where addrobj uses the naming format interface/user-defined-string, such as e1000g0/v4globalz.
```

Note:

```
create-addr [-t] -T static [-d] -a {local | remote}=addr[/prefixlen], ... addrobj
```

Creates a static IPv4 or IPv6 address on the interface specified in addrobj. If the interface on which the address is created is not plumbed, this subcommand will implicitly plumb the interface. The created static address will be identified by addrobj.

By default, a configured address will be marked up, so that it can be used as a source or destination of or for outbound and inbound packets.

### NEW QUESTION 10

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.

### NEW QUESTION 10

You log in to the system as user1, then switch user to root by using the su - command. After entering the correct password, you enter the following commands:

```
whoami;who am i;id
```

Which option correctly represents the output?

- A) uid=0(root) gid=0(root)  
user1 console Dec 30 20:20  
root
- B) root  
user1 console Dec 30 20:20  
uid=0(root) gid=0(root)
- C) user1 console Dec 30 20:20  
root  
uid=0(root) gid=0(root)
- D) uid=0(root) gid=0(root)  
root  
user1 console Dec 30 20:20

- A. Option A
- B. Option B
- C. Option C
- D. Option D

**Answer: B**

**Explanation:** \* The whoami utility displays your effective user ID as a name. Here this would be root.

\* who am i

The command who shows who is logged on. Here this would be:

user1 console Dec 30 20:20

\* The id utility displays the user and group names and numeric IDs, of the calling process, to the standard output. If the real and effective IDs are different, both are displayed, otherwise only the real ID is displayed.

Here this would be: uid=0(root) gid=0(root)

Note:

Each UNIX proces has 3 UIDs associated to it. Superuser privilege is UID=0.

Real UID

-----

This is the UID of the user/process that created THIS process. It can be changed only if the running process has EUID=0.

Effective UID

-----

This UID is used to evaluate privileges of the process to perform a particular action. EUID can be change either to RUID, or SUID if EUID!=0. If EUID=0, it can be changed to anything.

Saved UID

-----

If the binary image file, that was launched has a Set-UID bit on, SUID will be the UID of the owner of the file. Otherwise, SUID will be the RUID.

#### NEW QUESTION 11

The COMSTAR framework provides support for the iSCSI protocol. Select three options that correctly describe the COMSTAR framework.

- A. iSCSI devices can be used as dump devices.
- B. SCSI commands are carried over IP networks and enable you to mount disk devices from across the network onto your local system.
- C. Large amounts of data can be transferred over an IP network with very little network degradation.
- D. COMSTAR allows you to convert any Solaris11 host into a SCSI target device that can be accessed over a storage network.
- E. One IP port can handle multiple ISCSI target devices.

**Answer:** BDE

**Explanation:** B: By carrying SCSI commands over IP networks, the iSCSI protocol enables you to access block devices from across the network as if they were connected to the local system. COMSTAR provides an easier way to manage these iSCSI target devices.

D: Common Multiprotocol SCSI TARget, or COMSTAR, a software framework that enables you to convert any Oracle Solaris 11 host into a SCSI target device that can be accessed over a storage network by initiator hosts.

E: One IP port can handle multiple iSCSI target devices.

#### NEW QUESTION 15

You created an IP address for interface net3 with the following command, which executed successfully:

```
ipadm create-addr -T static -a 192.168.0.100/24 net3/v4
```

You then ran: ipadm show-if

The result indicated that the interface was down.

You then ran:

```
ipadm delete-addr net3/v4
```

```
ipadm create-addr -T static -a 192.168.0.101/24 net3/v4 ipadm show-if
```

The last command indicated that the interface was up.

Why did it work with the second address specified, but not the first?

- A. The 192.168.0.100 address is reserved for broadcast messages.
- B. Another device exists on the network, using the 192.168.0.100 address.
- C. The network interface card does not support the address 192.168.0.100.
- D. The address 192.168.0.100 is at a boundary and may not be configured in Oracle Solaris 11.
- E. 192.168.0.100 is a DHCP address and may not be statically configured in Oracle Solaris 11.

**Answer:** B

**Explanation:** The first IP address is already in use.

#### NEW QUESTION 17

User jack logs in to host Solaris and executes the following command sequence:

```
jack@solaris:~$ cd
jack@solaris:~$ ls -l testfile
-r-xrwxr-- 1 jack other 226 dec 20 20:20 testfile
jack@solaris:~$ id
uid=54326(jack) gid=1(other) groups=1(other)
jack@solaris:~$ id jill
uid=54327(jill) gid=1(other) groups=1(other)
```

Which three statements are correct?

- A. User jack can edit testfile because he has read and write permissions at the group level.
- B. User jack can use cat to output the contents of testfile because he has read permission as the file owner.
- C. User jill can change the permissions of testfile because she has write permission for the file at the group level.
- D. User jill can edit testfile because she has read and write permission at the group level.

- E. User jack can change permissions for testfile because he is the owner of the file.
- F. User jack can change permissions for testfile because he has execute permission for the file.

**Answer:** DEF

**NEW QUESTION 21**

You need to make sure that all of the software packages on your server are up to date. Without installing any updates, which two commands would display .my software updates that are available in the default Oracle repository?

- A. pkg list -u
- B. pkg verify -u '\*\*'
- C. pkg search -u
- D. pkg info -r '\*\*'
- E. pkg install -nv
- F. pkg update -nv '\*\*'

**Answer:** AD

**Explanation:** A: the pkg list command display a list of packages in the current image, including state and other information. By default, package variants for a different architecture or zone type are excluded.

D: pkginfo displays information about software packages that are installed on the system (with the first synopsis, with -l) or that reside on a particular device or directory (with the second synopsis, with -r).

Without options, pkginfo lists the primary category, package instance, and the names of all completely installed and partially installed packages. It displays one line for each package selected.

With -r, retrieve the data from the repositories of the image's configured publishers. Note that you must specify one or more package patterns in this case.

**NEW QUESTION 24**

Review the boot environment information displayed on your system:

```
oldBE      -    -    149.OK    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 two options accurately describe the newBE boot environment?

- A. It cannot be destroyed.
- B. It cannot be activated.
- C. It cannot be renamed.
- D. You can create a snapshot of it.
- E. It is activated but unbootable.
- F. It has been deleted and will be removed at the next reboot.

**Answer:** BC

**Explanation:** If the boot environment is unbootable, it is marked with an exclamation point (!) in the Active column in the beadm list output.

The beadm command restricts actions on unbootable boot environments as follows: You cannot activate an unbootable boot environment. (B)

You cannot destroy a boot environment that is both unbootable and marked as active on reboot.

You cannot create a snapshot of an unbootable boot environment.

You cannot use an unbootable boot environment or boot environment snapshot with the -e option of beadm create.

You cannot rename an unbootable boot environment. (C)

**NEW QUESTION 25**

The line

```
set noexec_user_stack=1
```

should be added to the /etc/system file to prevent an executable stack while executing user programs. What is the purpose of this?

- A. help prevent core dumps on program errors
- B. help programs to execute more quickly by keeping to their own memory space
- C. log any messages into the stack log
- D. help make buffer-overflow attacks more difficult

**Answer:** D

**Explanation:** How to Disable Programs From Using Executable Stacks Purpose: Prevent executable stack from overflowing. You must be in the root role.

Edit the /etc/system file, and add the following line: set noexec\_user\_stack=1

Reboot the system.

```
# reboot
```

**NEW QUESTION 28**

You have been tasked with creating a dedicated virtual network between two local zones within a single system, in order to isolate the network traffic from other zones on that system.

To accomplish this, you will create .

- A. an ether stub

- B. virtual router
- C. a virtual bridge
- D. a virtual network interface
- E. nothing, because a virtual switch is automatically created when the virtual network interfaces are created

**Answer:** D

**Explanation:** First create a virtual switch, then create a virtual network interface.

#### NEW QUESTION 31

When upgrading an existing system from Solaris 11 Express to Oracle Solaris 11, what happens to the datalink names?

- A. They follow the default naming convention for the newly installed version.
- B. They maintain their names.
- C. They are called eth#.
- D. They are called el00g#.
- E. They are left unnamed, to avoid conflicts, and need to be renamed after the installation process is complete.

**Answer:** A

**Explanation:** Network configuration in Oracle Solaris 11 includes

\* Generic datalink name assignment – Generic names are automatically assigned to datalinks using the net0, net1, netN naming convention, depending on the total number of network devices that are on the system

Note: There is no upgrade path from Oracle Solaris 10 to Oracle Solaris 11. You must perform a fresh installation.

#### NEW QUESTION 34

You want to install the openldap software package to a new boot environment for testing before introducing the new software package to the production environment. What option describes the correct procedure to:

- 1) create a new BE named newBE
- 2) install the software to that new BE only

- A. pkg install --newBE openldap
- B. pkg install --be-name newBE openldap
- C. beadm create newBE beadm mount newBE /mntpkg -R /mnt update openldap
- D. beadm create newBE beadm activate newBE pkg install openldap

**Answer:** D

**Explanation:** If you want to create a backup of an existing boot environment, for example, prior to modifying the original boot environment, you can use the beadm command to create and mount a new boot environment that is a clone of your active boot environment. This clone is listed as an alternate boot environment in the GRUB menu for x86 systems or in the boot menu for SPARC systems.

When you clone a boot environment by using the beadm create command, all supported zones in that boot environment are copied into the new boot environment.

How to Create a Boot Environment

1. Become the root role.
2. Create the boot environment.

```
# beadm create BeName
```

BeName is a variable for the name of the new boot environment. This new boot environment is inactive.

3. (Optional) Use the beadm mount command to mount the new boot environment.

```
# beadm mount BeName mount-point
```

Note: If the directory for the mount point does not exist, the beadm utility creates the directory, then mounts the boot environment on that directory.

If the boot environment is already mounted, the beadm mount command fails and does not remount the boot environment at the newly specified location.

4. (Optional) Activate the boot environment.

```
# beadm activate BeName
```

BeName is a variable for the name of the boot environment to be activated.

On reboot, the newly active boot environment is displayed as the default selection in the x86 GRUB menu or the SPARC boot menu.

#### NEW QUESTION 36

You have been asked to terminate a process that appears to be hung and will not terminate. The process table is shown below:

```
root 15163 15156 0 12:51:15 pts/3 0:00 hungscript What command will terminate the process?
```

- A. kill -9 15163
- B. kill -1 15163
- C. kill -15 15163
- D. kill -2 15163

**Answer:** A

**Explanation:** Here we should use SIGTERM to terminate the process. Note:

When no signal is included in the kill command-line syntax, the default signal that is used is

-15 (SIGKILL). Using the -9 signal (SIGTERM) with the kill command ensures that the process terminates promptly. However, the -9 signal should not be used to kill certain processes, such as a database process, or an LDAP server process. The result is that data might be lost.

Tip - When using the kill command to stop a process, first try using the command by itself, without including a signal option. Wait a few minutes to see if the process terminates before

using the kill command with the -9 signal.

#### NEW QUESTION 40

Which two accurately describe the Solaris IPS repository?

- A. It contains a collection of operating system patches.
- B. It contains a collection of software packages.
- C. All packages within an IPS package repository reside in a catalog.
- D. It is an ISO image of the Solaris installation media.
- E. The packages in a catalog are associated with a specific publisher.

**Answer:** BE

**Explanation:** Image Packaging System (IPS) is a new network based package management system included in Oracle Solaris 11. It provides a framework for complete software lifecycle management such as installation, upgrade and removal of software packages. IPS also enables you to create your own software packages, create and manage package repositories, and mirror existing package repositories.

Oracle Solaris software is distributed in IPS packages. IPS packages are stored in IPS package repositories, which are populated by IPS publishers.

E: The following command displays property information about the local repository.

```
$ pkgrepo get -s /export/repoSolaris11
```

```
SECTION PROPERTY VALUE publisher prefix solaris repository description This\ repository\ serves\ a\ copy\ of\ the\ Oracle\ Solaris\ 11\ Build\ 175b\ Package\
Repository. repository name Oracle\ Solaris\ 11\ Build\ 175b\ Package\ Repository
repository version 4
```

The value of the publisher prefix specifies that solaris is to be used in the following cases:

When more than one publisher's packages are present and no publisher is specified in the package name in the pkg command

When packages are published to the repository and no publisher is specified.

#### NEW QUESTION 42

View the Exhibit.



ADDR_OBJ	TYPE	STATE	ADDR
lo0/v4	static	ok	127.0.0.1/8
net0/_b	dhcp	ok	10.0.2.15/24
net1/_b	dhcp	ok	10.0.3.15/24
lo0/v6	static	ok	::1/128
net0/_a	addrconf	ok	fe80::a00:27ff:fee5:38b9/10
net1/_a	addrconf	ok	fe80::a00:27ff:fe2b:498a/10

After Installing the OS, you need to verify the network interface information. Which command was used to display the network interface information in the exhibit?

- A. ifconfig -a
- B. ipadm show-addr
- C. svcs -l network/physical
- D. netstat -a

**Answer:** B

**Explanation:** '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
```

#### NEW QUESTION 45

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.

**Answer:** C

**Explanation:** 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.

#### NEW QUESTION 47

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:** CE

**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

**NEW QUESTION 52**

View the Exhibit to inspect the boot environment Information displayed within a non global zone on your system.

BE/Dataset/Snapshot	Active	Mountpoint	Space	Policy	Created
solaris					
rpool/ROOT/solaris	NR	/	367.97M	static	2011-11-28 11:09
rpool/ROOT/solaris/var	-	-	26.16M	static	2011-11-28 11:09
rpool/ROOT/solaris/var@2011-11-28-18:49:38	-	-	69.0K	static	2011-11-28 13:49
rpool/ROOT/solaris/var@2011-11-28-19:09:23	-	-	0	static	2011-11-28 14:09
rpool/ROOT/solaris/var@install	-	-	975.0K	static	2011-11-28 12:29
rpool/ROOT/solaris@2011-11-28-18:49:38	-	-	70.0K	static	2011-11-28 13:49
rpool/ROOT/solaris@2011-11-28-19:09:23	-	-	0	static	2011-11-28 14:09
rpool/ROOT/solaris@install	-	-	929.5K	static	2011-11-28 12:29
solaris-1	!	-	2.0K	static	2011-11-28 13:49
rpool/ROOT/solaris-1	-	-	1.0K	static	2011-11-28 13:49
rpool/ROOT/solaris-1/var	-	-	57.0K	static	2011-11-28 14:09
z1BE					
rpool/ROOT/z1BE	-	-	1.0K	static	2011-11-28 14:09
rpool/ROOT/z1BE/var	-	-			

Which two options describe the solaris-1 boot environment?

- A. The solaris-1 boot environment is not bootable.
- B. The solaris-1 boot environment is incomplete.
- C. The solaris-1 boot environment was created automatically when the non global zone was created.
- D. The solaris-1 boot environment was created in the non-global zone using the beadm create command.
- E. The solaris-1 boot environment is associated with a non active global zone boot environment.

**Answer:** AE

**Explanation:** A: The ! of the Active Column indicates that this boot environment is inactive, and hence not bootable.

Note: The values for the Active column are as follows: R – Active on reboot.

N – Active now.

NR – Active now and active on reboot. “-” – Inactive.

“!” – Unbootable boot environments in a non-global zone are represented by an exclamation point.

[http://docs.oracle.com/cd/E23824\\_01/html/E21801/unbootable.html#scrolltoc](http://docs.oracle.com/cd/E23824_01/html/E21801/unbootable.html#scrolltoc)

**NEW QUESTION 53**

Which operation will fail if the DNS configuration is incorrect?

- A. domainname
- B. ping localhost.
- C. ping 192.168.1.1
- D. ping 23.45.82.174
- E. ping www.oracle.com.
- F. cat /etc/resolv.conf

**Answer:** E

**Explanation:** www.oracle.com would have to be resolved to an IP name by the domain name service.

**NEW QUESTION 56**

You start to execute a program by using the following command:

```
~/bigscript &
```

You then determine that the process is not behaving as expected, and decide that you need to terminate the process.

Based on the information shown below, what is the process number you should terminate?

```
#echo $$
15156
# ps -aef | grep 15156
  root 15163    15156    0   12:51:15   pts/3    0:00  bash
  root 15156     5420    0   12:33:15   pts/3    0:00  bash
  root 15166    15156    0   12:51:45   pts/3    0:00  grep
  root 15165    15156    0   12:51:45   pts/3    0:00  ps -aef
```

- A. 15163
- B. 15156
- C. 15166
- D. 15165

**Answer:** A

**Explanation:** From the output exhibit we can deduce that the shell has id 15156. It has spawned three subprocesses:

grep: id 15166

ps -aef 15165

The remaining 15163 must be the subshell (see note below). This is the id of the process which should be terminated.

#### NEW QUESTION 58

You are the administrator of a system that a large number of developers work on. These developers crash the system, and their applications, on a regular basis. What command would you use to configure where the core files are saved?

- A. savecore
- B. dumpadm
- C. svcadm
- D. proc
- E. coreadm

**Answer:** E

**Explanation:** The coreadm command is used to specify the name and location of core files produced by abnormally-terminating processes.

#### NEW QUESTION 60

You have a ZFS file system named /dbase/oral and you want to guarantee that 10 GB of storage space is available to that dataset for all data, snapshots, and clones.

Which option would you choose?

- A. zfs set refreservation=10g dbase/oral
- B. zfs set quota=10g dbase/oral
- C. zfs set refquota=10g dbase/oral
- D. zfs set reservation=10g dbase/oral

**Answer:** D

**Explanation:** A ZFS reservation is an allocation of disk space from the pool that is guaranteed to be available to a dataset. As such, you cannot reserve disk space for a dataset if that space is not currently available in the pool. The total amount of all outstanding, unconsumed reservations cannot exceed the amount of unused disk space in the pool. ZFS reservations can be set and displayed by using the zfs set and zfs get commands. For example:

```
# zfs set reservation=5G tank/home/bill
```

```
# zfs get reservation tank/home/bill NAME PROPERTY VALUE SOURCE
```

```
tank/home/bill reservation 5G local
```

#### NEW QUESTION 61

Which two statements describe the COMSTAR framework available in Oracle Solaris 11?

- A. It converts an Oracle Solaris 11 host into a SCSI target device that can be accessed over a storage network by Linux, Mac OS, or Windows client systems.
- B. iSCSI targets cannot be configured as dump devices.
- C. It provides support for iSCSI devices that use SLP.
- D. It is used to connect to Fibre Channel or iSCSI Storage Area Network (SAN) environments.
- E. It provides an upgrade and update path to convert your iSCSI LUNs from Solaris 10 systems.

**Answer:** AB

**Explanation:** A: You can configure Common Multiprotocol SCSI TARget, or COMSTAR, a software framework that enables you to convert any Oracle Solaris 11 host into a SCSI target device that can be accessed over a storage network by initiator hosts. This means you can make storage devices on a system available to Linux, Mac OS, or Windows client systems as if they were local storage devices. Supported storage protocols are iSCSI, FC, iSER, and SRP.

B: iSCSI targets cannot be configured as dump devices.

#### NEW QUESTION 63

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
```

#### NEW QUESTION 65

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: ADE**

**Explanation:** A: Snapshots are created by using the zfs snapshot command, which takes as its only argument the name of the snapshot to create.

D: You can list snapshots as follows:

```
# zfs list -t snapshot
```

E: Snapshots are destroyed by using the zfs destroy command. For example:

```
# zfs destroy tank/home/ahrens@now
```

#### NEW QUESTION 68

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

- 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: AC**

**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.

#### NEW QUESTION 69

The core dump configuration for your system is:

```
global core file pattern: /var/core/core.%f.%p
global core file content: default
init core file pattern: core.%f.%p.%z
init core file content: default
global core dumps: enabled
per-process core dumps: enabled
global setid core dumps: enabled
per-process setid core dumps: enabled
global core dump logging: disabled
```

A user is running a process in the global zone and the process crashes. The process information is:

```
User1 2663 2618 0 17:46:42 pts/2 0:00 /usr/bin/bash
```

The server host name is: zeus

What will the per-process core file be named?

- A. core.bash.2663.global
- B. core.bash.2663.zeus
- C. /var/core/core.bash.2663
- D. /var/core/core.bash.2663.global

**Answer: C**

**Explanation:** Note the first line:

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

The program name is bash The runtime process ID is 2663

Note: By default, the global core dump is disabled. You need to use the coreadm command with the -e global option to enable it. The -g option causes the command to append the program name(%f) and the runtime process ID (%p) to the core file name.

#### NEW QUESTION 71

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

```
pool: pool1
state: ONLINE
scan: none requested
config:
      NAME                STATE          READ  WRITE  CKSUM
      pool1                ONLINE         0     0     0
      mirror-0             ONLINE         0     0     0
      c3t3d0                ONLINE         0     0     0
      c3t4d0                ONLINE         0     0     0
```

The server has two spare 146-GB disk drives: c3t5d0 c3t6d0

You need to add more space to the pool1 storage pool. Which command would add more mirrored storage to the pool1 storage pool?

- A. zpool add pool1 mirror c3t5d0 c3t6d0
- B. zpool attach pool1 mirror c3t5d0 c3t6d0
- C. zpool attach pool1 c3r3d0 c3r5d0; zpool attach pool1 c3r4d0 c3r6d0
- D. zpool add pool1 c3r3d0 c3r5d0; zpool add pool1 c3r4d0 c3r6d0

**Answer:** A

#### NEW QUESTION 74

In order to display the IP addresses of network interfaces, what command would you use?

- A. dladm
- B. ipconfig
- C. sves
- D. ipadm
- E. ipaddr

**Answer:** D

**Explanation:** 'ipadm show-addr' displays all the configured addresses on the system. Example:

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

#### NEW QUESTION 77

User jack logs in to host solar in and issues the following command:

```
jack@solaris:~$ ls .ssh
```

id\_dsa id\_dsa.pub id\_rsa id\_rsa.pub known\_hosts authorized\_keys Which two are true?

- A. The id\_rsa file contains the private key for rhosts-based host authentication.
- B. The id\_dsa.pub file contains the Digital Signature Algorithm public key for the user jack.
- C. The id\_rsa.pub file contains the Rivest Shamir Adelman public key for the host solaris.
- D. The authorized\_keys file contains the private keys of remote users authorized to access jack's account on solaris.
- E. The known\_hosts file contains the verified public keys of remote hosts known to be trusted.

**Answer:** AE

**Explanation:** A: You will see two files starting with id\_rsa. id\_rsa is the private key and id\_rsa.pub is public key.

E: The .ssh/known\_hosts file

In order to use public-key secure connection with other hosts (ssh, scp, sftp) there is a special directory, ~/.ssh/, where passphrases and public keys are stored. Normally you wouldn't need to know the gory details, but from time to time a host will change its public key and then you have difficulty using ssh or scp with that host, and have to edit a file named known\_hosts.

If you try to ssh to another computer, but get an error message that warns about a changed or incorrect public key, then it is probably just a case of that host changing its public key. (It is possible, though usually not the case, that malicious hacking is involved.) Unless you actually suspect hacker involvement, you can edit the file ~/.ssh/known\_hosts using your usual text editor (vi, emacs, nedit, or pico) and delete any line with the name of that host.

Then when you try to ssh that host again, it will be like the first time ever; ssh will ask you if you want to accept a new public key, you type the whole word yes, and everything will proceed normally from there.

Here is what a typical ~/.ssh/known\_hosts file might contain. Note that newton is represented on two different lines:

```
newton 1024 35
153438062610297067329638677441205712613292203533062535600064224677647442
245028855505387934431717435134842994423656065076260604296084868001730665
553662299156116414854701274715680961503198280525759778667306417179500370
```

```
189017139564144825610347509023078143132936185076849630461827976942220442
313116255293297021841
ucsub 1024 37
132170811640421742212085598383135714069016332111955003414250071326834884
018721183646445780180633494496866895830879394309011412231102757022090299
732775466435482517698989962531081214859205054227533597152962802400251809
883548442498002326460312850336779152617243800769119880843882425555806081
435017335194477605333
simpson 1024 41
840896920592494584403453622735282634536002054701576247765078766974814128
393752943151071629834843909016027026612791643752972116459602750267266908
365259665072736159491719667576217171370458928680504368847255632477925660
234893185547218857655484574619075125368470792976275806263534208879722192
77539015703446529603
newton, 128.138.249.8 ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEA0d7Aoure0toNJ+YMYi61QP2ka8m5x5ZQIT7obP8C
K3eropfqsMPPY6uiyh9vpiFX2r1LHcbx139+vG6HOtVvuS8+IfMDtawm3WQvRuOopz3vVy
5GtMwtaOgehsXoT930Ryev1bH5myPtWKlipITsOd2sX9k3tvjrmme4KCGGss=
```

**NEW QUESTION 82**

In a default standalone installation of Oracle Solaris 11, what is the default minimum length in characters of a user password, and where is the minimum password length defined?

- A. Default minimum length is 8, and is defined in /etc/default/password.
- B. Default minimum length is 6, and is defined in /etc/default/password.
- C. Default minimum length is 8, and is defined in /etc/shadow.
- D. Default minimum length is 6, and is defined in /etc/shadow.
- E. Default minimum length is 8, and is defined in /usr/sadm/defadduser.
- F. Default minimum length is 6, and is defined in /usr/sadm/defadduser.

**Answer: B**

**Explanation:** By default, the passwd command assumes a minimum length of six characters. You can use the PASSLENGTH default in the /etc/default/passwd files to change that by setting the minimum number of characters that a user's password must contain to some other number.

**NEW QUESTION 86**

Identify the correct description of an IPS image.

- A. An ISO image of the Solaris media DVD
- B. An IPS repository
- C. A depot location or source where Solaris packages can be installed from
- D. A location where packages can be installed, for example, your Solaris instance

**Answer: D**

**Explanation:** An image is a location where packages can be installed. An image can be one of three types:

- \* Full images are capable of providing a complete system.
- \* Partial images are linked to a full image (the parent image), but do not provide a complete system on their own.
- \* User images contain only relocatable packages.

**NEW QUESTION 87**

When issuing the zonestat 2 1h is command, the following information is displayed:

```
SUMMARY          Cpus/Online: 1/1   PhysMem: 1023M  VirtMem: 2047M
---CPU---      --PhysMem--  --VirtMem--  --PhysNet--
ZONE  USED  %PART  USED  %USED  USED  %USED  PBYTE  %PUSE
[total] 0.09 9.33%  841M 82.1%  951M 46.4%  0 0.00%
[system] 0.02 2.40%  319M 31.2%  577M 28.1%  -  -
global 0.06 6.71%  465M 45.4%  325M 15.8%  0 0.00%
dbzone 0.00 0.21%  56.1M 5.48%  48.7M 2.37%  0 0.00%
```

Which two options accurately describe the statistics contained in the output?

- A. dbzone is using 0.21% of the total CPU resource available in the zone's processor set.
- B. dbzone is using 0.21% of the global zone's total CPU.
- C. dbzone is using 5.48% of the total physical memory that has been allocated to the zone.
- D. dbzone is using 2.37% of the global zone's total virtual memory.
- E. The network is being utilized 100% with no physical bandwidth remaining.

**Answer: AC**

**Explanation:** A: %PART

The amount of cpu used as a percentage of the total cpu in a processor-set to which the zone is bound. A zone can only have processes bound to multiple processor sets if it is the global zone, or if psrset(1m) psets are used. If multiple binding are found for a zone, it's %PART is the fraction used of all bound psets. For [total] and [system], %PART is the percent used of all cpus on the system.

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.

**NEW QUESTION 88**

Which option displays the result of running the zfs list command?

- A) 

NAME	SIZE	ALLOC	FREE	CAP	DEDUP	HEALTH	ALTROOT
pool1	15.9G	144K	15.9G	0%	1.00x	ONLINE	-
- B) 

NAME	USED	AVAIL	REFER	MOUNTPOINT
pool1	144K	15.6G	31K	none
- C) 

```
pool: pool1
state: ONLINE
scan: none requested
config:
      NAME      STATE      READ WRITE CKSUM
      pool1     ONLINE     0     0     0
      c3t3d0    ONLINE     0     0     0
```
- D) 

pool	capacity		operations		bandwidth	
	alloc	free	read	write	read	write
pool1	144K	15.9G	0	0	62	754
rpool	6.35G	9.52G	5	1	44.4K	10.6K
zone	3.41G	12.5G	0	0	76	17

- A. Option A
- B. Option B
- C. Option C
- D. Option D

**Answer: B**

**Explanation:** The zfs list command provides an extensible mechanism for viewing and querying dataset information.

You can list basic dataset information by using the zfs list command with no options. This command displays the names of all datasets on the system and the values of their used, available, referenced, and mountpoint properties. For more information about these properties, see *Introducing ZFS Properties*.

For example:

```
# zfs list
NAME USED AVAIL REFER MOUNTPOINT
pool 476K 16.5G 21K /pool
pool/clone 18K 16.5G 18K /pool/clone pool/home 296K 16.5G 19K /pool/home
pool/home/marks 277K 16.5G 277K /pool/home/marks pool/home/marks@snap 0 - 277K -
pool/test 18K 16.5G 18K /test
```

**NEW QUESTION 89**

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

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

**Answer: B**

**Explanation:** 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.

**NEW QUESTION 91**

Which two capabilities are provided by the OpenBoot PROM?

- A. a command to safely shut down the system
- B. hardware testing and initialization
- C. booting from a disk or network
- D. starting the GRUB loader

**Answer: BC**

**Explanation:** OpenBoot firmware is executed immediately after you turn on your system. The primary tasks of OpenBoot firmware are to:

- \* Test and initialize the system hardware (B)
- \* Determine the hardware configuration
- \* Boot the operating system from either a mass storage device or from a network (C)
- \* Provide interactive debugging facilities for testing hardware and software

#### NEW QUESTION 92

You have installed an update to the gzip package and need to "undo" .ho update and return the package to its "as-delivered" condition. Which command would you use?

- A. pkg undo
- B. pkg revert
- C. pkg fix
- D. pkg uninstall

**Answer: B**

**Explanation:** Use the pkg revert command to restore files to their as-delivered condition.

#### NEW QUESTION 93

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.

#### NEW QUESTION 96

Oracle Solaris 11 kernel encounters a fatal error, and it results in a system panic. What type of file does this generate?

- A. a.out
- B. objdump
- C. core dump
- D. tape dump
- E. crash dump

**Answer: C**

**Explanation:** A kernel panic is a type of error that occurs when the core (kernel) of an operating system receives an instruction in an unexpected format or when it fails to handle properly. A kernel panic can also follow when the operating system can't recover from a different type of error. A kernel panic can be caused by damaged or incompatible software or, more rarely, damaged or incompatible hardware.

When a server kernel panics it abruptly halts all normal system operations. Usually, a kernel process named panic() outputs an error message to the console and stores debugging information in nonvolatile memory to be written to a crash log file upon restarting the computer. Saving the memory contents of the core and associated debugging information is called a "core dump."

#### NEW QUESTION 101

Which three of the components could be used in a ZFS storage pool, but are not recommended configurations?

- A. A file on a UFS file system
- B. A Veritas Volume Manager (VxVM) volume
- C. A LUN in a hardware RAID array
- D. A disk slice from an SMI labeled disk
- E. A Solaris Volume Manager (SVM) volume
- F. An EFI labeled disk

**Answer: ABE**

**Explanation:** A: ZFS also allows you to use UFS files as virtual devices in your storage pool. This feature is aimed primarily at testing and enabling simple experimentation, not for production use. The reason is that any use of files relies on the underlying file system for consistency. If you create a ZFS pool backed by files on a UFS file system, then you are implicitly relying on UFS to guarantee correctness and synchronous semantics.

However, files can be quite useful when you are first trying out ZFS or experimenting with more complicated layouts when not enough physical devices are present. All files must be specified as complete paths and must be at least 64 Mbytes in size.

B, E: You can construct logical devices for ZFS using volumes presented by software-based volume managers, such as Solaris Volume Manager (SVM) or Veritas Volume Manager (VxVM). However, these configurations are not recommended. While ZFS functions properly on such devices, less-than-optimal performance might be the result.

#### NEW QUESTION 106

Identify the Automated Installer's (AI) equivalent to jumpStart's finish scripts and sysidcfg files.

- A. Manifest files

- B. SMF system configuration profile files
- C. Installadm create - client
- D. IPS software package repository
- E. installadm create-service
- F. svccfg - s application/pkg/server setprop sysidcfg

**Answer:** B

**Explanation:** Comparing sysidcfg File Keywords to System Configuration Profile Directives

The following table compares sysidcfg file keywords with example AI system configuration profile specifications.

sysidcfg File Keyword

System Configuration Profile Directives Etc.

#### NEW QUESTION 108

Identify three differences between the shutdown and init commands.

- A. Only shutdown broadcasts a final shutdown warning to all logged-in users.
- B. init does not terminate all services normal
- C. The shutdown command performs a cleaner shutdown of all services.
- D. The shutdown command can only bring the system to the single-user milestone
- E. The init command must be used to shut the system down to run level 0.
- F. Only shutdown sends a shutdown message to any systems that are mounting resources from the system that is being shut down.
- G. The shutdown command will shut the system down and turn off power; init will only shut the system down.

**Answer:** ABE

#### NEW QUESTION 111

You are installing the Solaris 11 OE by using the Interactive Text Installer. You have selected the option to automatically configure the primary network controller. Which three items will automatically be configured as a result of this selection?

- A. The IP address.
- B. The name service.
- C. The time zone.
- D. A default user account.
- E. The terminal type.
- F. The root password.
- G. The host name.

**Answer:** ABC

**Explanation:** IP address and name service (such as a DNS server) are provided by the DHCP server.

#### NEW QUESTION 116

You enter `dladm show-phys`, which provides the following output:

LINK	MEDIA	STATE	SPEED	DUPLEX	DEVICE
net0	ethernet	up	1000	full	e1000g1
net3	ethernet	up	1000	full	e1000g3

You then enter: `ipadm create-ip net3`

What is the output?

- A. `ipadm: cannot; create interface net3: Operation failed.`
- B. `ipadm: cannot create interface net3: Interface already exists.`
- C. `ipadm: cannot create interface net3: IP address object not specified.`
- D. `No_response, The command was successful.`

**Answer:** B

**Explanation:** According to the exhibit the interface already exists.

The command `ipadm create-ip net3` is supposed to create a new interface net3.

#### NEW QUESTION 118

You are logged in as root to a newly installed Solaris 11 system. You issue the command `useradd -d`, and then examine the `/usr/sadm/defadduser` file. This file includes the entry `defshell=/bin/sh`. Which shell will now be the default for the next account created?

- A. bash shell
- B. C shell
- C. korn shod
- D. bourne shell

**Answer:** A

**Explanation:** Oracle Solaris 11 introduces user environment and command-line argument changes that include the following:

\* Shell changes - The default shell, `/bin/sh`, is now linked to `ksh93`. The default user shell is the Bourne-again (bash) shell.

\* The legacy Bourne shell is available as `/usr/sunos/bin/sh`.

- \* The legacy ksh88 is available as /usr/sunos/bin/ksh from the shell/ksh88 package.
- \* Korn shell compatibility information is available in /usr/share/doc/ksh/COMPATIBILITY.

### NEW QUESTION 123

To confirm the IP address and netmask have been correctly configured on the network interfaces which command should you use?

- A. ipdilm show-if
- B. ipadm show-nic
- C. ipadm show-addr
- D. ipadm show-ifconfig
- E. ipadm show-addr ipadm show-mask

**Answer: C**

**Explanation:** Show address information, either for the given addrobj or all the address objects configured on the specified interface, including the address objects that are only in the persistent configuration.

State can be: disabled, down, duplicate, inaccessible, ok, tentative Example:

```
# ipadm show-addr
```

```
ADDROBJ TYPE STATE ADDR
```

```
lo0/v4 static ok 127.0.0.1/8 lo0/v6 static ok ::1/128
```

### NEW QUESTION 127

You need to migrate a UFS file system named /production\_ufs to a ZFS file system named /production\_zfs. The /production\_ufs file system cannot be taken down or be out of production during the migration, and the current /production\_ufs file system must remain active until the /production\_zfs file system is copied and ready.

Which method allows you to meet both requirements?

1. Copy live data from /production\_ufs to /production\_zfs while /production\_ufs is in use.
2. When the copy is complete, /production\_zfs will contain an up-to date copy of /production\_ufs

- A. Create a snapshot of the UFS file system
- B. Create the new ZFS file system
- C. Use cpio to copy data from the snapshot to the new ZFS file system.
- D. Create a new Boot Environment
- E. Create the ZFS file system
- F. Use lucreate -m to copy data from the Current UFS file system to the new ZFS file system.
- G. Mirror the existing UFS file system by using SVM. After both submissions are in sync, migrate one of the submissions to a ZFS file System by using Live Upgrade.
- H. Create the new ZFS file system by using zfs create import to import data from the existing UFS file system into the new ZFS file system
- I. Create the new zfs file system by using the zfs create -o shadow.

**Answer: E**

**Explanation:** Migrating Data With ZFS Shadow Migration

ZFS shadow migration is a tool you can use to migrate data from an existing file system to a new file system. A shadow file system is created that pulls data from the original source as necessary.

You can use the shadow migration feature to migrate file systems as follows:

- \* A local or remote ZFS file system to a target ZFS file system
- \* A local or remote UFS file system to a target ZFS file system

Shadow migration is a process that pulls the data to be migrated:

- \* Create an empty ZFS file system.
- \* Set the shadow property on an empty ZFS file system, which is the target (or shadow) file system, to point to the file system to be migrated.

For example:

```
# zfs create -o shadow=nfs://system/export/home/ufsdata users/home/shadow2
```

\* Data from file system to be migrated is copied over to the shadow file system.

### NEW QUESTION 130

You need to install the gzip software package on your system. Which command would you use to find the software package in the configured repository?

- A. pkg search gzip
- B. pkg info gzip
- C. pkg contents gzip
- D. pkginfo gzip
- E. yum list gzip

**Answer: A**

**Explanation:** Use the pkg search command to search for packages whose data matches the specified pattern.

Like the pkg contents command, the pkg search command examines the contents of packages. While the pkg contents command returns the contents, the pkg search command returns the names of packages that match the query.

### NEW QUESTION 131

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:** BEF

**Explanation:** B: The mount point Error. The following message appears during the boot process or in response to an explicit mount request and indicates a non-existent mount point.

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.

#### NEW QUESTION 134

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
- F. The image must be the same version as the Oracle Solaris OS that you plan to install on the client.
- G. Download the text install image into the IPS repository.
- H. Install the AI installation tools.
- I. Create the AI install service
- J. Specify the path to the AI network boot image ISO file and the path where the AI net image ISO file should be unpacked.
- K. Create the AI install service
- L. Specify the path to the AI network boot image ISO file and the path to the IPS repository.

**Answer:** BDFGI

**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.

#### NEW QUESTION 135

The following information is displayed for the svc:/network/ssh service:

```

fmri          svc:/network/ssh:default
name          SSH server
enabled       true
state         offline
next_state    none
state_time    December 31, 2011 07:10:08 AM EST
logfile       /var/svc/log/network-ssh:default.log
restarter     svc:/system/svc/restarter:default
contract_id   321
manifest      /etc/svc/profile/generic.xml
manifest      /lib/svc/manifest/network/ssh.xml
dependency    require_all/none svc:/system/filesystem/local (online)
dependency    optional_all/none svc:/system/filesystem/autofs (online)
dependency    require_all/none svc:/network/loopback (online)
dependency    require_all/none svc:/network/physical:default (online)
dependency    require_all/none svc:/system/cryptosvc (disabled)
dependency    require_all/none svc:/system/utmp (online)
dependency    optional_all/error svc:/network/ipfilter:default (disabled)
dependency    require_all/restart file:///localhost/etc/ssh/sshd_config (online)

```

```

svc:/network/ssh:default (SSH server)
State: offline since January 31, 2012 09:12:45 AM EST
Reason: Service svc:/system/cryptosvc:default is disabled.
  See: http://sun.com/msg/SMF-8000-GE
  Path: svc:/network/ssh:default
        svc:/system/cryptosvc:default
  See: man -M /usr/share/man -s 1M sshd
  See: /var/svc/log/network-ssh:default.log
Impact: This service is not running.

```

Which describes the minimum set of commands to be executed to bring the svc:/network/ssh: default service back online?

- A) `svcadm refresh svc:/network/ssh:default`
- B) `svcadm restart svc:/network/ssh:default`
- C) `svcadm enable svc:/system/cryptosvc`
- D) `svcadm enable svc:/system/cryptosvc`  
`svcadm enable svc:/network/ipfilter:default`  
`svcadm enable svc:/network/ssh:default`
- E) `svcadm enable svc:/system/cryptosvc`  
`svcadm enable svc:/network/ipfilter:default`  
`svcadm refresh svc:/network/ssh:default`
- F) `svcadm restart svc:/system/cryptosvc`  
`svcadm restart svc:/network/ipfilter:default`  
`svcadm restart svc:/network/ssh:default`
- G) `svcadm enable svc:/network/ssh:default`

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E
- F. Option F
- G. Option G

Answer: C

#### NEW QUESTION 139

Given the following output of the zpool status command:

```
pool: pool1
state: ONLINE
scan: none requested
config:
  NAME          STATE      READ      WRITE     CKSUM
  pool1         ONLINE    0         0         0
  raidz1-0     ONLINE    0         0         0
  c3t3d0       ONLINE    0         0         0
  c3t4d0       ONLINE    0         0         0
  c3t5d0       ONLINE    0         0         0
  c3t6d0       ONLINE    0         0         0
```

Identify the correct statement regarding pool1's configuration.

- A. Data written to pool1 will be striped across four disk components.
- B. The rsdz1-0 and c3t640 components are submirrors of pool1.
- C. Data will only be striped across the three disks in rsidz configuration.
- D. The configuration is a bug in Solaris 11; it cannot be created by an administrator.

Answer: B

#### NEW QUESTION 140

ServerA contains two ISO images of a package repository named so1.repo.iso-a and so1.repo.iso-b respectively. You need to create a single local package repository on server that clients can connect to. The package repository will be stored on the /export/IPS file system and named repo. The preferred publisher will be named solaris and the publisher URL will be http://serverA.example.com.

Which is the correct procedure to perform on ServerA to create the local Package repository?

- A. cat so1.repo.iso-a sol.repo.iso-b > so1.full.isoMount the ISO image and use the rsync command to extract the contents of the ISO file to the /export/IPS file system.Set the pkg/inst\_root property to /export/IPS/repo and the pkg/readonly property to true.Set the preferred publisher by using pkg set-publisher -Ghttp://pkg.oracle.com/solaris/release/ \-g http://serverA.example.com/ solaris
- B. cat so1.repo.iso-a so1.repo.iso-b > /export/IPS/repoSet the pkg/inst\_root property to true and the pkg/readonly property to /export/IPSSet the preferred publisher by using pkg set-publisher -G http://serverA.example.com/ \-g http://pkg/oracle.com/solaris/rekease/solaris
- C. cat so1.repo.iso-a so1.repo.iso-b > so1.full.isoMount the ISO image and use the rsync command to extract the contents of the ISO file to /export/IPS/repoSet the pkg/inst\_root property to /export/IPS/repo and the pkg/readonly property to trueSet the preferred publisher by using pkg set-publisher solaris \-g http://pkg.oracle.com/
- D. cat so1.repo, iso-a so1.repo.iso-b > /export/IPS/repo.isoMount the ISO image and copy the repo directory from the ISO image to /export/IPS/repoSet the pkg/inst\_root property and the pkg/readonly property to /export/IPS/repoSet the preferred pkg/inst\_root property by using pkg set-publisher - G http://serverA.example.com/ \- g http://pkg.oracle.com/solaris.com/release/- p solaris

Answer: A

#### NEW QUESTION 141

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

- A. If a location is explicitly enabled, it remains active until explicitly changed.
- B. Wireless security keys can be configured by using the nwammgr command.
- C. NWAM stores profile information in /etc/ipadm/ipadm.conf and /etc/dladm/datalink.conf.
- D. Multiple locations may be automatically activated in systems with multiple network interface cards.
- E. Interface NCU Properties "float" and are automatically attached to the highest priority Link NCU Property.
- F. If the DefaultFixed NCP is enabled, persistent configuration, stored in /etc/ipadm.conf and /etc/dladm/datalink.conf is used.

Answer: AD

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

D: A location comprises certain elements of a network configuration, for example a name service and firewall settings, that are applied together, when required. You can create multiple locations for various uses. For example, one location can be used when you are connected at the office by using the company intranet. Another location can be used at home when you are connected to the public Internet by using a wireless access point. Locations can be activated manually or automatically, according to environmental conditions, such as the IP address that is obtained by a network connection.



#### NEW QUESTION 144

You have Solaris 11 system with a host name of sysA and it uses LDAP as a naming service. You have created a flash archive of sysA and you want to migrate this system to an Oracle Solaris11 server, Solaris10 branded zone. The zone Status on the Oracle Solaris 11 server is:  
 - zone10 incomplete/zone/zone1solaris10exc1  
 Select the option that will force the non-global zone to prompt you for a host name and name service the first time it is booted.

- A. Use zonecfg to change the zonename before booting the system for the first time
- B. Use the -u option with the zoneadm -z zone10 attach command.
- C. Use the -u option with the zoneadm -z zone10 install command.
- D. Remove the sysidcfg file from the <zonepath>/root directory before booting the non- global zone.

**Answer: C**

**Explanation:** Oracle Solaris 10 branded zones – Oracle Solaris 10 Zones provide an Oracle Solaris 10 environment on Oracle Solaris 11. You can migrate an Oracle Solaris 10 system or zone to a solaris10 zone on an Oracle Solaris 11 system in the following ways:

- \* Create a zone archive and use the archive to create an s10zone on the Oracle Solaris 11 system. This option applies in the current scenario.
- Example of command to Install the Oracle Solaris 10 non-global zone. s11sysB# zoneadm -z s10zone install -u -a /pond/s10archive/s10.flar
- \* Detach the zone from the Oracle Solaris 10 system and attach the zone on the Oracle Solaris 11 zone. The zone is halted and detached from its current host. The zonepath is moved to the target host, where it is attached.
- Note:  
 install [-x nodataset] [brand-specific options] A subcommand of the zoneadm.  
 Install the specified zone on the system. This subcommand automatically attempts to verify first. It refuses to install if the verify step fails.  
 -u uuid-match  
 Unique identifier for a zone, as assigned by libuuid(3LIB). If this option is present and the argument is a non-empty string, then the zone matching the UUID is selected instead of the one named by the -z option, if such a zone is present.

#### NEW QUESTION 149

User jack on host solaris attempts to use ssh to log in to host oracle and receives this message:  
 jack@solaris:~\$ ssh oracle  
 ssh: connect to host oracle port 22: connection refused What is the problem?

- A. Host oracle does not have a valid host public key.
- B. Host oracle does not have a valid host private key.
- C. Host solaris does not have a valid host public key.
- D. Host does not have a valid host private key.
- E. Host solaris is not configured for host-based authentication.
- F. Host oracle is not configured for host-based authentication.
- G. Host oracle is not running the ssh service.
- H. Host solaris is not running the ssh service.

**Answer: G**

**Explanation:** The host he is trying to connect to (oracle) is not running the required service (ssh).

#### NEW QUESTION 150

You need to update an OS image on a client. The pkg publishers command displays the wrong publisher with the wrong update:  
 PUBLISHERTYPESTATUSURI  
 Solaris origin onlinehttp://pkg.oracle.com/solaris/release  
 The update is available on the updated publisher: PUBLISHERTYPESTATUSURI  
 Solaris originonlinehttp://sysA.example.com

Select the option that describes the procedure used to update the OS image on the system from the updated publisher.

- A. Copy the repository from the ISO image onto the local client
- B. Configure the repository on the client by using the `svccfg - s` command so that the Solaris publisher is connected to the new repository
- C. Refresh the application/pkg/server service
- D. Issue the `pkgrepo refresh` command to refresh the repository catalog
- E. Configure the publisher on the client using the `svcfg - s` command so that the Solaris publisher is connected to the repository at `http://sysA.example.com` Refresh the application/pkg/server service
- F. Issue the `pkgrepo refresh` command to repository catalog
- G. Use the `pkg set-publisher` command to change the URL of the publisher Solaris to `http://sysA.example.co`
- H. Issue the `pkg update` command to update the OS image.
- I. Add the new publisher `http://sysA.example.com` Solaris Use the `pkg set-publisher` command to set the publisher search order and place `http://sysA.example.com` of `http://pkg.oracle.com/solaris/release` Issue the `pkg publisher` command to view the publisher
- J. Set the new publisher to stick
- K. Issue the `pkg update` command to update the OS image.

**Answer: C**

**Explanation:** You can use the `pkg set-publisher` command to change a publisher URI. Changing a Publisher Origin URI

To change the origin URI for a publisher, add the new URI and remove the old URI. Use the `-g` option to add a new origin URI. Use the `-G` option to remove the old origin URI.

```
# pkg set-publisher -g http://pkg.example.com/support \
-G http://pkg.example.com/release example.com
```

Note: You can use either the `install` or `update` subcommand to update a package.

The `install` subcommand installs the package if the package is not already installed in the image. If you want to be sure to update only packages that are already installed, and not install any new packages, then use the `update` subcommand.

#### NEW QUESTION 153

You need to configure an iSCSI target device on your x86 based Oracle Solaris II system. While configuring the iSCSI device, the following error is displayed:

```
bash: stmfadm: command not found
```

Which option describes the solution to the problem?

- A. The COMSTAR feature is not supported on the x86 platform
- B. The feature is supported only on the SPARC platform.
- C. Use the `iscsitadm` command on the x86 platform when configuring an iSCSI target.
- D. Install the storage-server group package on this system.
- E. Start the iSCSI target daemon on this system.

**Answer: C**

**Explanation:** STMF – Manages transactions, such as context and resources for Small Computer System Interface (SCSI) command execution, and tracking logical unit and port providers. STMF also handles logical unit mappings, allocating memory, recovering failed operations, enumeration, and other necessary functions of an I/O stack.

STMF is controlled by `stmfadm`, and `stmfadm` is the majority of the commands you will be using to administer COMSTAR (COMmon Multiprotocol Scsi TARget).

Install the packages you need for COMSTAR with iSCSI and reboot:

```
# pfexec pkg install storage-server
# pfexec pkg install SUNWiscsit
# shutdown -y -i6 -g0
```

Note: You can set up and configure a COMSTAR Internet SCSI (iSCSI) target and make it available over the network. The iSCSI features can work over a normal Internet connection (such as Ethernet) using the standard iSCSI protocol. The iSCSI protocol also provides naming and discovery services, authentication services using CHAP and RADIUS, and centralized management through iSNS.

The COMSTAR target mode framework runs as the `stmf` service. By default, the service is disabled. You must enable the service to use COMSTAR functionality.

You can identify the service with the `svcs` command. If you have not rebooted the server since installing the group/feature/storage-server package, the service might not be enabled correctly.

#### NEW QUESTION 155

You need to install the `solaris-desktop` group package. Which command would you use to list the set of packages included in that software group?

- A. `pkg search`
- B. `pkg info`
- C. `pkginfo`
- D. `pkg contents`

**Answer: A**

**Explanation:** Use the `pkg search` command to search for packages whose data matches the specified pattern.

Like the `pkg contents` command, the `pkg search` command examines the contents of packages. While the `pkg contents` command returns the contents, the `pkg search`

command returns the names of packages that match the query.

#### NEW QUESTION 157

You have been asked to do an orderly shutdown on a process with a PID of 1234, with the `kill` command.

Which command is best?

- A. `kill -2 1234`
- B. `kill -15 1234`
- C. `kill -9 1234`

D. kill -1 1234

**Answer:** B

**Explanation:** On POSIX-compliant platforms, SIGTERM is the signal sent to a process to request its termination. The symbolic constant for SIGTERM is defined in the header file signal.h. Symbolic signal names are used because signal numbers can vary across platforms, however on the vast majority of systems, SIGTERM is signal #15.

SIGTERM is the default signal sent to a process by the kill or killall commands. It causes the termination of a process, but unlike the SIGKILL signal, it can be caught and interpreted (or ignored) by the process. Therefore, SIGTERM is akin to asking a process to terminate nicely, allowing cleanup and closure of files. For this reason, on many Unix systems during shutdown, init issues SIGTERM to all processes that are not essential to powering off, waits a few seconds, and then issues SIGKILL to forcibly terminate any such processes that remain.

#### NEW QUESTION 159

The storage pool configuration on your server is:

```
pool1          200K      3.91G   31K    /pool1
pool1/data     31K       3.91G   31K    /pool1/data
pool1          ONLINE     0       0       0
c4t0d0        ONLINE     0       0       0
```

You back up the /pool1/data file system, creating a snapshot and copying that snapshot to tape (/dev/rmt/0). You perform a full backup on Sunday night and incremental backups on Monday through Saturday night at 11:00 pm. Each incremental backup will copy only the data that has been modified since the Sunday backup was started.

On Thursday, at 10:00 am, you had a disk failure. You replaced the disk drive (c4t0d0). You created pool (pool1) on that disk.

Which option would you select to restore the data in the /pool1/data file system?

- A. zfs create pool1/dataLoad the Monday tape and enter:zfs recv pool1/data </dev/rmt/0Load the Wednesday tape and enter:zfs recv -F pool1/data < /dev/rmt/0
- B. Load the Sunday tape and restore the Sunday snapshot:zfs recv pool1/data </dev/rmt/0zfs rollback pool1/data@monLoad the Wednesday tape and restore the Wednesday snapshot:zfs recv -i pool1/data < /dev/rmt/0zfs rollback pool1/data@wed
- C. zfs create pool1/dataLoad the Wednesday tape and enter:zfs recv -F pool1/data </dev/rmt/0
- D. Load the Sunday tape and enter:zfs recv pool1/data < /dev/rmt/0Load the Wednesday tape and enter:\* commands missing\*

**Answer:** D

**Explanation:** First the full backup must be restored. This would be the Sunday backup.

Then the last incremental backup must be restored. This would be the Wednesday backup. Before restoring the Wednesday incremental file system snapshot, the most recent snapshot must first be rolled back.

By exclusion D) would be best answer even though it is incomplete.

#### NEW QUESTION 163

To assist in examining and debugging running processes, Solaris 11 has a utility that returns pro arguments and the names and values of environment variables. What is the name of this utility?

- A. ppgsz
- B. pargs
- C. pmap
- D. pgrep

**Answer:** B

**Explanation:** The pargs utility examines a target process or process core file and prints arguments, environment variables and values, or the process auxiliary vector.

#### NEW QUESTION 168

You have installed software updates to a new boot environment (BE) and have activated that the booting to the new BE, you notice system errors. You want to boot to the last known good configuration.

Which option would you use on a SPARC system to boot to the currentBE boot environment?

- A. boot -L currentBE
- B. boot -Z rpool/ROOT/currentBE
- C. boot -a Enter the currentBE dataset name when prompted.
- D. boot rpool/ROOT/currentBE
- E. boot -m currentBE
- F. beadm activate currentBE

**Answer:** F

**Explanation:** You can change an inactive boot environment into an active boot environment. Only one boot environment can be active at a time. The newly activated boot environment becomes the default environment upon reboot.

How to Activate an Existing Boot Environment

1. Use the following command to activate an existing, inactive boot environment: beadm activate beName

beName is a variable for the name of the boot environment to be activated. Note the following specifications.

beadm activate beName activates a boot environment by setting the bootable pool property, bootfs, to the value of the ROOT dataset of the boot environment that is being activated.

beadm activate sets the newly activated boot environment as the default in the menu.lst file.

2. Reboot.

The newly activated boot environment is now the default on the x86 GRUB menu or SPARC boot menu.

**NEW QUESTION 171**

Your are troubleshooting network throughput on your server.

To confirm that the load balancing among aggregated links is functioning properly, you want to examine the traffic statistics on the links comprising the aggregation.

The correct command is .

- A. dlstat - aggr
- B. dlstat show-aggr
- C. dlstat show-link -r
- D. dlstat show-link -aggr
- E. dlstat show-phys -aggr

**Answer: B**

**Explanation:** dlstat show-aggr [-r | -t] [-i interval] [-p] [-o field[, ...]] [-u R|K|M|G|T|P] [link] Display per-port statistics for an aggregation.

**NEW QUESTION 175**

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.

**Answer: E**

**Explanation:** 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
```

**NEW QUESTION 180**

View the Exhibit and review the disk configuration.

```
Specify disk (enter its number)[2]:
selecting c3t3d0
[disk formatted]
format> ver

Primary label contents:

Volume name = <          >
ascii name = <ATA-VBOX HARDDISK-1.0 cyl 2085 alt 2 hd 255 sec 63>
pcyl      = 2087
ncyl      = 2085
acyl      = 2
bcyl      = 0
nhead     = 255
nsect     = 63
Part      Tag      Flag      Cylinders      Size      Blocks
0 unassigned  wm        0 - 130        1.00GB      (131/0/0) 2104515
1 unassigned  wm        0                0          (0/0/0)    0
2 backup      wu        0 - 2084       15.97GB     (2085/0/0) 33495525
3 unassigned  wm        0                0          (0/0/0)    0
4 unassigned  wm        0                0          (0/0/0)    0
5 unassigned  wm        0                0          (0/0/0)    0
6 unassigned  wm        0                0          (0/0/0)    0
7 unassigned  wm        0                0          (0/0/0)    0
8 boot       wu        0 - 0          7.84MB     (1/0/0)    16065
9 unassigned  wm        0                0          (0/0/0)    0

format> █
```

The following command is executed on the disk: zpool create pool1 c3t3d0s0

What is the result of executing this command?

- A. A zpool create error is generated.
- B. A 1-GB ZFS file system named /pool1 is created.
- C. A 15.97-GB storage pool named pool1 is created.
- D. The disk will contain an EFI disk label.

Answer: B

**NEW QUESTION 183**

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.

**NEW QUESTION 185**

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 endBoot the zone and log in the zone as root
- B. Create the zpool: zpool create pool2 c4t0d0 In the non-global zone, root can now create ZFS file system in the pool2 zpool
- C. 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 endadd fsset dir=pool1 set special=pool1 set type=zfs pool1 endBoot the zone, log in the zone as root, and create the zfs file system in the pool2 zpool.
- D. 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 endBoot the zone, log in to the zone as root and create the zfs file systems in the pool2 zpool.
- E. 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
- F. 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

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

**NEW QUESTION 188**

Your SPARC server will not boot into multi user-server milestones and you need to troubleshoot to out why. You need to start the server with minimal services running so that you can go through each milestone manually to troubleshoot the issue.

Select the option that boots the server with the fewest services running.

- A. boot -s

- B. boot milestone none
- C. boot -m milestone=single-user
- D. boot -m milestone=none
- E. boot -m none

**Answer:** D

**Explanation:** The command boot -m milestone=none is useful in repairing a system that have problems booting early.

Boot Troubleshooting:

To step through the SMF portion of the boot process, start with: boot -m milestone=none

Then step through the milestones for the different boot levels: svcadm milestone svc:/milestone/single-user:default

svcadm milestone svc:/milestone/multi-user:default svcadm milestone svc:/milestone/multi-user-server:default

#### NEW QUESTION 189

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:** BC

**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.

#### NEW QUESTION 191

You have a ticket from a new user on the system, indicating that he cannot log in to his account. The information in the ticket gives you both the username and password. The ticket also shows that the account was set up three days ago.

As root, you switch users to this account with the following command: su - newuser

You do not get an error message.

You then run ls -la and see the following files:

```
local1.cshrc local1.login local1.profile .bash_history .bashrc .profile
```

As root, you grep the /etc/passwd file and the /etc/shadow file for this username, with these results:

```
/etc/passwd contains newuser:x:60012:10:/home/newuser:/usr/bin/bash
```

```
/etc/shadow contains newuser:UP: : : :10: : As root, what is your next logical step?
```

- A. Usermod -f 0
- B. passwd newuser
- C. mkdir /home/newuser
- D. useradd -D

**Answer:** B

**Explanation:** The content of the /etc/shadow document indicates that the newuser account has no password.

We need to add a password.

The passwd utility is used to update user's authentication token(s). D: Here the user account already exist. There is no need to create it.

When invoked without the -D option, the useradd command creates a new user account using the values specified on the command line plus the default values from the system. Depending on command line options, the useradd command will update system files and may also create the new user's home directory and copy initial files.

#### NEW QUESTION 196

A user jack, using a bash shell, requests a directory listing as follows:

```
jack@solaris:~$ ls
dira dirb dirc diraa dirabc
```

Which three statements are correct?

- A. The pattern dir? will expand to dira dirb dirc.
- B. The pattern dir\*a will expand to diraa.
- C. The pattern dir\*a will expand to dira diraa.
- D. The pattern dir\*b? will expand to dirabc.
- E. The pattern dir\*b? will expand to dirb dirabc.

**Answer:** ACD

**Explanation:** A: dir followed by a single letter.

- C: dir followed by any characters ending with a.
- D: dir followed by any characters, then character b, then one single character. only dirabc matches

#### NEW QUESTION 201

Which statement is correct about shutdown and init commands?

- A. shutdown broadcasts one or more periodic shutdown warning messages to all logged-in users whereas init issues none.
- B. The shutdown command performs a clean shutdown of all services whereas init does not.
- C. The shutdown command brings the system to the single-user milestone by default
- D. The init command must be used to shut the system down to run level 0.
- E. The shutdown command accepts SMF milestones, init stages, or run levels as arguments whereas init accepts only init stages or run levels as arguments.

**Answer:** A

#### NEW QUESTION 205

Which two are implemented using the Internet Control Message Protocol (ICMP)?

- A. ping
- B. DHCP
- C. HTTP
- D. telnet
- E. syslog
- F. traceroute

**Answer:** AF

**Explanation:** The Internet Control Message Protocol (ICMP) is one of the core protocols of the Internet Protocol Suite. ICMP differs from transport protocols such as TCP and UDP in that it is not typically used to exchange data between systems, nor is it regularly employed by end-user network applications (with the exception of some diagnostic tools like ping and traceroute).

#### NEW QUESTION 209

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?

- 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/dhccpinfo DNSserv to ensure that the DHCP server providing the DNS server information.

**Answer:** ABF

**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/dhccpinfo 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.

#### NEW QUESTION 213

Identify the two security features incorporated in the Oracle Solaris 11 Cryptographic Framework.

- A. Layer 5 IP address encryptions
- B. Internet protocol security
- C. Diffie-Kerberos coaxial key encryption
- D. Signed cryptographic plugins (providers)
- E. Kernel support for signed antivirus plugins

**Answer:** DE

**Explanation:** The framework enables providers of cryptographic services to have their services used by many consumers in the Oracle Solaris operating system. Another name for providers is plugins. The framework allows three types of plugins:

\* User-level plugins - Shared objects that provide services by using PKCS #11 libraries, such as pkcs11\_softtoken.so.1.

\* Kernel-level plugins - Kernel modules that provide implementations of cryptographic algorithms in software, such as AES.

Many of the algorithms in the framework are optimized for x86 with the SSE2 instruction set and for SPARC hardware.

\* Hardware plugins - Device drivers and their associated hardware accelerators. The Niagara chips, the ncp and n2cp device drivers, are one example. A hardware accelerator offloads expensive cryptographic functions from the operating system. The Sun Crypto Accelerator 6000 board is one example.

#### NEW QUESTION 218

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.

**Answer: B**

**Explanation:** 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.

SU 03/31 12:52 + pts/0 <userid>-root

### NEW QUESTION 222

When you issue the “gzip: zommand 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**

**Explanation:** Searching for Packages

Use the pkg search command to search for packages whose data matches the specified pattern.

Like the pkg contents command, the pkg search command examines the contents of packages. While the pkg contents command returns the contents, the pkg search command returns the names of packages that match the query.

pkg search

search [-Hlflpr] [-o attribute ...] [-s repo\_uri] query Search for matches to the query, and display the results.

Which tokens are indexed are action-dependent, but may include content hashes and pathnames.

Note: pkg is the retrieval client for the image packaging system. With a valid configuration, pkg can be invoked to create locations for packages to be installed, called 'images', and install packages into those images. Packages are published by publishers, who may make their packages available at one or more repositories. pkg, then, retrieves packages from a publisher's repository and installs them into an image.

### NEW QUESTION 223

You have installed the SMF notification framework to monitor services. Which command is used to set up the notifications for a particular service?

- A. svccfg
- B. svcadm
- C. setnotify
- D. smtp-notify

**Answer: A**

**Explanation:** How to Set Up Email Notification of SMF Transition Events

This procedure causes the system to generate an email notification each time one of the services or a selected service has a change in state. You can choose to use either SMTP or SNMP. Normally, you would only select SNMP if you already have SNMP configured for some other reason.

By default, SNMP traps are sent on maintenance transitions. If you use SNMP for monitoring, you can configure additional traps for other state transitions.

1. Become an administrator or assume a role that includes the Service Management rights profile.'

2. Set notification parameters. Example 1:

The following command creates a notification that sends email when transactions go into the maintenance state.

```
# /usr/sbin/svccfg setnotify -g maintenance mailto:sysadmins@example.com
```

Example 2:

The following command creates a notification that sends email when the switch service goes into the online state.

```
# /usr/sbin/svccfg -s svc:/system/name-service/switch:default setnotify to-online \ mailto:sysadmins@example.com
```

Note: The svccfg command manipulates data in the service configuration repository. svccfg can be invoked interactively, with an individual subcommand, or by specifying a command file that contains a series of subcommands.

Changes made to an existing service in the repository typically do not take effect for that service until the next time the service instance is refreshed.

### NEW QUESTION 226

Your server has one zone named dbzone (hat has been configured, but not yet installed). Which command would you use to view all the options that were used to configure this zone?

- A. zoneadm list -icv dbzone
- B. zones tat -c summary dbzone
- C. zonecfg -z dbzone info
- D. zonecfg -icv dbzone info

**Answer: C**

**Explanation:** zonecfg info

Display information about the current configuration. If resource-type is specified, displays only information about resources of the relevant type. If any property-name value pairs are specified, displays only information about resources meeting the given criteria. In the resource scope, any arguments are ignored, and info displays information about the resource which is currently being added or modified.

Note: zonecfg -z

zonename. Specify the name of a zone. Zone names are case sensitive. Zone names must begin with an alphanumeric character and can contain alphanumeric characters, the underscore (\_) the hyphen (-), and the dot (.). The name global and all names beginning with SUNW are reserved and cannot be used.

Incorrect Answer

A: The zoneadm utility is used to administer system zones. A zone is an application container that is maintained by the operating system runtime.

list option:

Display the name of the current zones, or the specified zone if indicated. B: No such command.

D: no such options zonecfg -icv

**NEW QUESTION 228**

You are attempting to edit your crontab file in the bash shell. Instead of getting your usual vi interface, you are presented with an unfamiliar interface. In order to have your editor of choice-vi- what command must you type after exiting the unfamiliar editor?

- A. EDITOR=vi
- B. crontab=vi
- C. crontab - e vi
- D. env

**Answer:** A

**Explanation:** Set the EDITOR variable to vi.

Commands like `crontab -e` will use ed per default. If you'd like to use some better editor (like vi) you can use the environment variable EDITOR:

# EDITOR=vi; crontab -e will open the users crontab in vi. Of course you can set this variable permanently.

Incorrect answers

C: -e Edits a copy of the current user's crontab file, or creates an empty file to edit if crontab does not exist. When editing is complete, the file is installed as the user's crontab file. If a username is given, the specified user's crontab file is edited, rather than the current user's crontab file; this can only be done by a user with the solaris.jobs.admin authorization. The environment variable EDITOR determines which editor is invoked with the -e option. The default editor is ed(1). All crontab jobs should be submitted using crontab. Do not add jobs by just editing the crontab file, because cron is not aware of changes made this way.

**NEW QUESTION 232**

The following line is from /etc/shadow in a default Solaris 11 Installation:

jack: \$5\$9JFrt54\$7JdwmO.F11Zt/ jFeeOhDmnw93LG7Gwd3Nd/cwCcNWFFg:0:15:30:3::: Which two are true?

- A. Passwords for account jack must be a minimum of 15 characters long.
- B. The password for account jack has expired.
- C. The password for account jack has 5 characters.
- D. A history of 3 prior passwords for the account jack is kept to inhibit password reuse.
- E. The minimum lifetime for a password for account jack is 15 days.

**Answer:** BE

**Explanation:** From the content of the /etc/shadow file we get:

\* username: jack

\* encrypted password: \$5\$9JFrt54\$7JdwmO.F11Zt/ jFeeOhDmnw93LG7Gwd3Nd/cwCcNWFFg

\* Last password change (lastchanged): Days since Jan 1, 1970 that password was last changed: 0

\* Minimum: The minimum number of days required between password changes i.e. the number of days left before the user is allowed to change his/her password: 15

Maximum: The maximum number of days the password is valid (after that user is forced to change his/her password): 30 Warn : The number of days before password is to expire that user is warned that his/her password must be changed: 3

\* Inactive : The number of days after password expires that account is disabled

\* Expire : days since Jan 1, 1970 that account is disabled i.e. an absolute date specifying when the login may no longer be used

**NEW QUESTION 237**

Which command would you use to determine which package group is installed on your system?

- A. pkg list group/system/\*
- B. pkg info
- C. uname -a
- D. cat /var/sadm/system/admin/CLUSTEP

**Answer:** B

**Explanation:** The pkg info command provides detailed information about a particular IPS package. Note: The pkginfo command does the same for any SVR4 packages you may have installed on the same system.

pkg info example:

\$ pkg info p7zip Name: compress/p7zip

Summary: The p7zip compression and archiving utility

Description: P7zip is a unix port of the 7-Zip utility. It has support for numerous compression algorithms, including LZMA and LZMA2, as well as for various archive and compression file formats, including 7z, xz, bzip2, gzip, tar, zip (read-write) and cab, cpio, deb, lz, rar, and rpm (read-only).

Category: System/Core State: Installed Publisher: solaris Version: 9.20.1

Build Release: 5.11

Branch: 0.175.0.0.0.2.537

Packaging Date: Wed Oct 19 09:13:22 2011  
 Size: 6.73 MB  
 FMRI: pkg://solaris/compress/p7zip@9.20.1, 5.11-0.175.0.0.0.2.537:20111019T091322Z

**NEW QUESTION 239**

Which three statements are true concerning Image Packaging System (IPS) incorporation package?

- A. Installing an incorporation package does not install any other packages.
- B. Every feature or tool has a separate IPS incorporation.
- C. They constrain the versions of packages they incorporate.
- D. They are a content management tool and not a version management tool.
- E. Their dependencies are always of TYPE-REQUIRE.
- F. They are defined by their manifest

**Answer:** ACE

**NEW QUESTION 241**

You have set up the task.max-lwps resource control on your Solaris 11 system.  
 Which option describes how to configure the system so that syslogd notifies you when the resources control threshold value for the task.max-lwps resource has been exceeded?

- A. Use the rctldm command to enable the global action on the task.max-lwpa resource control.
- B. Modify the /etc/syslog.conf file to activate system logging of all violations of task.max-lwps and then refresh then svc: /system/system-log:default service.
- C. Activate system logging of all violations of task.max-lwpp in the /etc/rctldm.conf file and then execute the rctldm-u command.
- D. Use the prct1 command to set the logging of all resource control violations at the time the task.max-lwps resource control is being setup.
- E. Use the setrct1 command to set the logging of all resource control violations for the task.max-lwps resource control.

**Answer:** A

**Explanation:** rctldm - display and/or modify global state of system resource controls

The following command activates system logging of all violations of task.max-lwps.

```
# rctldm -e syslog task.max-lwps
#
```

**NEW QUESTION 244**

After installing the OS, the following network configuration information is displayed from the system:

```
ADDBOBJ      TYPE      STATE      ADDR
lo0/v4       static    ok         127-0.0.1/8
lo0/v6       static    ok         ::1/128
```

Which option describes the state of this server?

- A. The automatic network configuration option was chosen during the installation of the OS.
- B. The manual network configuration option was chosen during the installation of the OS.
- C. The network was not configured during the installation of the OS.
- D. The network interface is configured with a static IP address.

**Answer:** C

**Explanation:** Only the loopback addresses are configured. No IP address is configured.

**NEW QUESTION 245**

On localSYS, your SPARC based server, you back up the root file system with recursive snapshots of the root pool. The snapshots are stored on a remote NTS file system.

This information describes the remote system where the snapshots are stored:

Remote system name: backupSYS

File system where the snapshots are stored: /backups/localSYS Mounted file system on localSYS: /rpool/snaps

Most recent backup name: rpool-1202

Disk c0t0d0 has failed in your root pool and has been replaced. The disk has already been

part< and labeled and now you need to restore the root file system. Which procedure would you follow to restore the ZFS root file system on localSYS?

- A. boot cdrom -smount -f nfs backup\_server:/rpool/snaps /rmtzpool create rpool c0t0d0s0cat /mnt/rpool.1202 | zfs receive -Fdu rpoolzpool set bootfs=rpool/ROOT/solaris rpoolRecreate swap and dump devices.Reinstall the bootblock on c0t0d0.
- B. boot cdrom -smount -f nfs backup\_server:/rpool/snaps /mntzpool create rpool c0t0d0s0zfs create -o mountpoint=/ rpool/ROOTcat /mnt/rpool.1011 | zfs receive -Fdu rpoolzpool set bootfs=rpool/ROOT/solaris rpoolRecreate swap and dump devices.Reinstall the bootblock on c0t0d0.
- C. boot cdrom -smount -F nfs backup\_server:/rpool/snaps /mntcat /mnt/rpool.1011 | zfs receive -Fdu rpoolzpool set bootfs=rpool/ROOT/solaris rpool c0t0d0s0Reinstall the bootblock on c0t0d0s0
- D. boot cdrom -smount -f nfs backup\_server:/rpool/snaps /rmtzpool create rpool c0t0d0s0zfs receive -Fdu /mnt/rpool.1011zpool set bootfs=rpool/ROOT/solaris rpoolReinstall the bootblock on c0t0d0.

**Answer:** A

**Explanation:** How to Recreate a ZFS Root Pool and Restore Root Pool Snapshots In this scenario, assume the following conditions:

\* ZFS root pool cannot be recovered

\* ZFS root pool snapshots are stored on a remote system and are shared over NFS

\* The system is booted from an equivalent Solaris release to the root pool version so that the Solaris release and the pool version match. Otherwise, you will need to add the -o version=version-number property option and value when you recreate the root pool in step 4 below.

All steps below are performed on the local system. 1.

Boot from CD/DVD or the network.

On a SPARC based system, select one of the following boot methods:

ok boot net -s

ok boot cdrom -s

If you don't use -s option, you'll need to exit the installation program.

2.

Mount the remote snapshot dataset. For example:

```
# mount -F nfs remote-system:/rpool/snaps /mnt
```

3.

Recreate the root pool. For example:

```
# zpool create -f -o failmode=continue -R /a -m legacy -o cachefile=/etc/zfs/zpool.cache rpool c1t0d0s0
```

4.

Restore the root pool snapshots.

This step might take some time. For example:

```
# cat /mnt/rpool.0311 | zfs receive -Fdu rpool
```

Using the -u option means that the restored archive is not mounted when the zfs receive operation completes.

5.

Set the bootfs property on the root pool BE. For example:

```
# zpool set bootfs=rpool/ROOT/osalBE rpool 6.
```

Install the boot blocks on the new disk.

On a SPARC based system:

```
# installboot -F zfs /usr/platform/`uname -i`/lib/fs/zfs/bootblk /dev/rdisk/c1t0d0s0
```

### NEW QUESTION 247

You want to display the IP address assignments of the network interfaces. Which command should you use?

- A. ipadm show-if
- B. ipadm show-addr
- C. ipadm show-prop
- D. ipadm show-addrprop

**Answer: B**

**Explanation:** 'ipadm show-addr' displays all the configured addresses on the system. Example:

```
# ipadm show-addr
```

```
ADDROBJ TYPE STATE ADDR
```

```
lo0/v4 static ok 127.0.0.1/8 lo0/v6 static ok ::1/128
```

### NEW QUESTION 252

zone1 is a non-global zone that has been configured and installed.

zone1 was taken down for maintenance, and the following command was run: zoneadm -z zone1 mark incomplete

The following information is displayed when listing the zones on your system:

ID	NAME	STATUS	PATH	BRAND	IP
0	global	running	/	solaris	shared
-	dbzone	installed	/export/dbzone	solaris	excl
-	zone1	incomplete	/zone/zone1	solaris10	excl

Which task needs to be performed before you can boot zone1?

- A. The zone needs to be installed.
- B. The zone needs to be brought to the ready state.
- C. The zone needs to be uninstalled and reinstalled.
- D. The zone needs to be brought to the complete state.

**Answer: C**

**Explanation:** If administrative changes on the system have rendered a zone unusable or inconsistent, it is possible to change the state of an installed zone to incomplete.

Marking a zone incomplete is irreversible. The only action that can be taken on a zone marked incomplete is to uninstall the zone and return it to the configured state.

### NEW QUESTION 255

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?

- 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
```

**NEW QUESTION 259**

You need to know the IP address configured on interface net3, and that the interface is up. Which command confirms these?

- A. ipadm show-if
- B. ipadm up-addr
- C. ipadm show-addr
- D. ipadm enable-if
- E. ipadm refresh-addr
- F. ipadm show-addrprop

**Answer:** C

**Explanation:** Show address information, either for the given addrobj or all the address objects configured on the specified interface, including the address objects that are only in the persistent configuration.

State can be: disabled, down, duplicate, inaccessible, ok, tentative Example:

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

**NEW QUESTION 263**

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

**Answer:** C

**Explanation:** 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.

**NEW QUESTION 267**

View the Exhibit to see the information taken from the installation log file.

Based on the information presented in the Exhibit, which two options describe the state of the system when the server is booted for the first time after the installation is complete?

```
Disk: 16.0GB sata
Partition: 16.0GB Solaris2
Slice 0: 16.0GB rpool

Time Zone: US/Eastern

Language: *The following can be changed when logging in.
Default language: English

Keyboard layout: *The following can be changed when logging in.
Default keyboard layout: US-English

Terminal type: sun-color

Users:
No user account

Network:
Computer name: solaris
Network Configuration: None
```

- A. NWAM will be used to configure the network interface.
- B. The network/physical service is offline.
- C. You cannot log in from the console as roo
- D. You must first log in as a user and then su to root account.
- E. The root user can log in from the console login.
- F. You will be prompted to configure the network interface after the initial login.

**Answer:** BD

#### NEW QUESTION 268

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?

- 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:** ABCF

**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.

#### NEW QUESTION 270

You have a user that needs to use the cron tool to schedule some repetitive tasks. When the user enters the crontab -e command in a terminal window, the following error appears:

```
crontab: you are not authorized to use cron. Sorry
```

In order to troubleshoot this issue, in what directory would you start your invest

- A. /etc/cron.d
- B. /var/spool/cron
- C. /var/spool/cron/crontable
- D. /var/spool/cron/atjobs

**Answer:** A

**Explanation:** crontab: you are not authorized to use cron. Sorry.

This message means that either the user is not listed in the cron.allow file (if the file exists), or the user is listed in the cron.deny file.

You can control access to the crontab command by using two files in the /etc/cron.d directory: cron.deny and cron.allow. These files permit only specified users to perform crontab command tasks such as creating, editing, displaying, or removing their own crontab files.

The cron.deny and cron.allow files consist of a list of user names, one user name per line.

#### NEW QUESTION 272

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **1Z0-821 Practice Exam Features:**

- \* 1Z0-821 Questions and Answers Updated Frequently
- \* 1Z0-821 Practice Questions Verified by Expert Senior Certified Staff
- \* 1Z0-821 Most Realistic Questions that Guarantee you a Pass on Your First Try
- \* 1Z0-821 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The 1Z0-821 Practice Test Here](#)**