

1Z0-821 Dumps

Oracle Solaris 11 System Administrator

<https://www.certleader.com/1Z0-821-dumps.html>



NEW QUESTION 1

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.

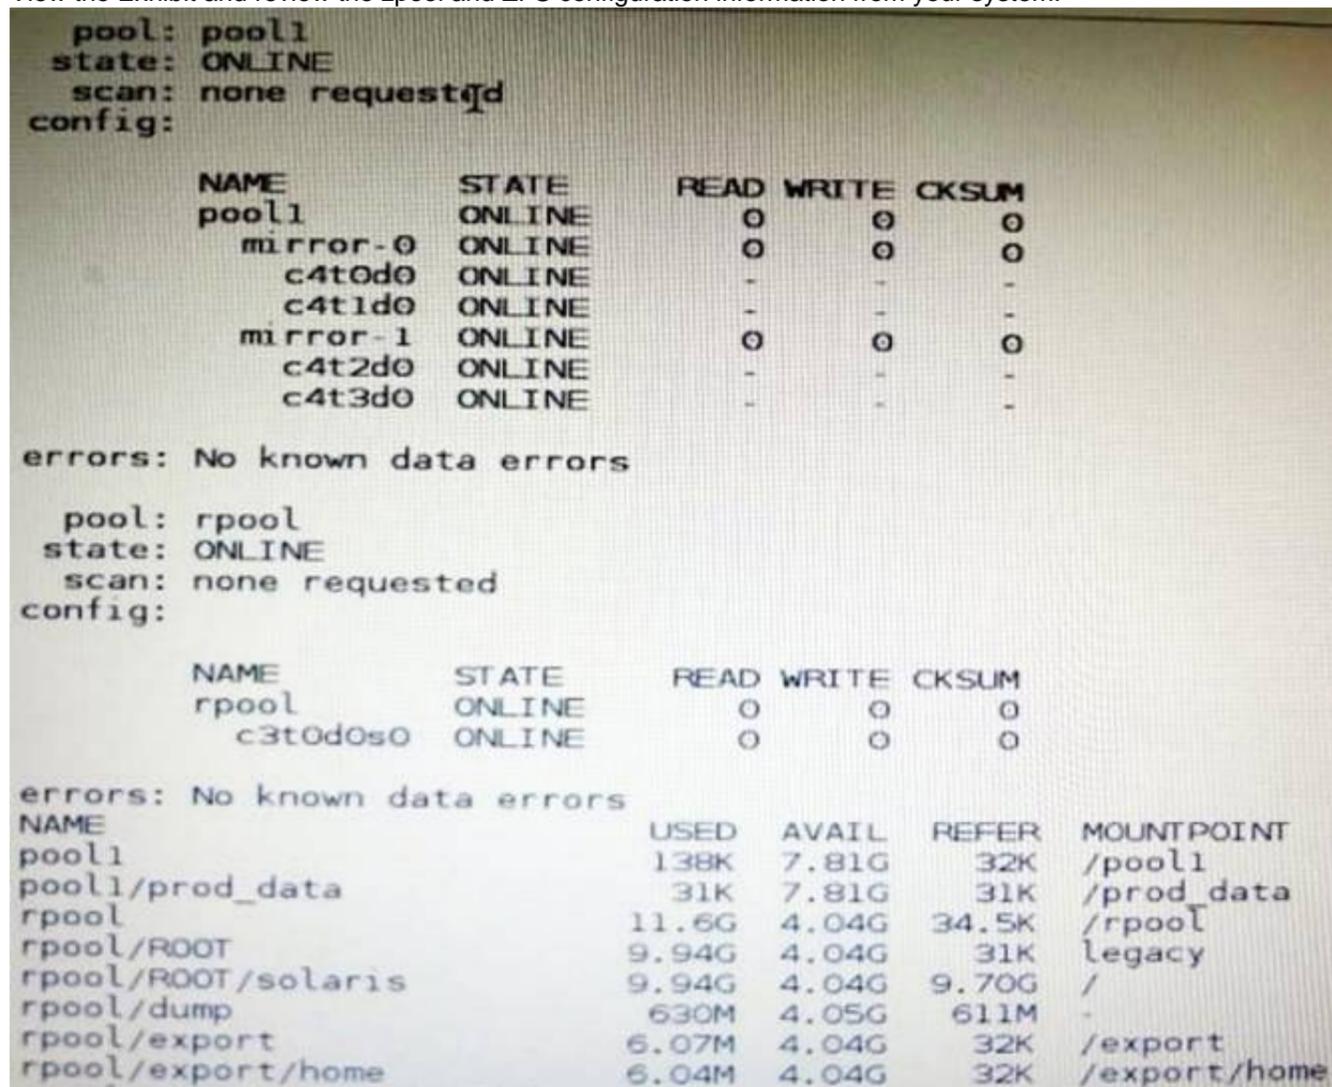
beadm list

BE Active Mountpoint Space Policy Created

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

NEW QUESTION 2

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



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 3

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 4

You are installing Oracle Solaris 11 on a SPARC-based system by using the Test Installer. Which three statements are true?

- A. The ROOT user will always be configured as a role.
- B. The root filesystem will always be deployed on ZFS.
- C. The root filesystem will always be located on a local disk.
- D. The network can be configured using DHCP.
- E. The set of packages that will be installed are server based.
- F. You must always create one regular user when installing the system.

Answer: BDE

NEW QUESTION 5

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 6

You suspect a problem with the openldap package and want to make sure that the files have not be modified or otherwise tampered with.

Which command would validate all of the files contained in the openldap package and report any problems?

- A. pkgchk openldap
- B. pkginfo openldap
- C. pkg contents openldap
- D. pkg verify openldap
- E. pkg set-property signature-policy verify

Answer: A

Explanation:

pkgchk checks the accuracy of installed files or, by using the -l option, displays information about package files. pkgchk checks the integrity of directory structures and files. Discrepancies are written to standard error along with a detailed explanation of the problem.

NEW QUESTION 7

You are going to create live zones on you server. Disk space is critical on this server so you need to reduce the amount of disk space required for these zones.

Much of the data required for each of these zones is identical, so you want to eliminate the duplicate copies of data and store only data that is unique to each zone.

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

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

Answer: DE

Explanation:

In Oracle Solaris 11, you can use the deduplication (dedup) property to remove redundant data from your ZFS file systems. If a file system has the dedup property enabled, duplicate data blocks are removed synchronously. The result is that only unique data is stored, and common components are shared between files.

NEW QUESTION 8

The crash dump notification on your server is:

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

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

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

Which command should you choose to display these two files?

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

Answer: A

Explanation:

Decompress using savecore -f vmdump.0

savecore - save a crash dump of the operating system

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

NEW QUESTION 9

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 10

What determines which bits in an IP address represent the subnet, and which represent the host?

- A. Subnet
- B. unicast
- C. netmask
- D. multicast
- E. broadcast

Answer: C

Explanation:

A subnetwork, or subnet, is a logically visible subdivision of an IP network. The practice of dividing a network into two or more networks is called subnetting.

The routing prefix of an address is written in a form identical to that of the address itself. This is called the network mask, or netmask, of the address. For example, a specification of the most-significant 18 bits of an IPv4 address, 11111111.11111111.11000000.00000000, is written as 255.255.192.0.

NEW QUESTION 10

Which command would you use from the bash shell to determine the total amount of physical memory installed in your Solaris system (x86 and SPARC)?

- A. uname -a
- B. prtconf | grep -i memory
- C. sysdef | grep -i memory
- D. vmstat

E. `prtdiag | grep -i memory`

Answer: B

Explanation:

The `prtconf` command prints the system configuration information. The output includes the total amount of memory, and the configuration of system peripherals formatted as a device tree.

If a device path is specified on the command line for those command options that can take a device path, `prtconf` will only display information for that device node.

NEW QUESTION 13

User brian changes the permissions for `db_data` this command: `chmod 4755 db_data`

What is true?

- A. `db_data` now has permissions `rwsr-xr-x` and can be deleted only by user brian.
- B. `db_data` now has permissions `rwsr-xr-x` and, if executed, will run with the permissions of user brian.
- C. `db_data` now has permissions `rw-r-xr-x` and can be deleted only by members of the group owning it.
- D. The permissions for `db_data` cannot be determined, because the permissions prior to the change have not been specified.
- E. `db_data` must be an ordinary file, because special permissions cannot be set on a directory.

Answer: C

Explanation:

Use the `chmod` command to change permissions for a file or directory. You must be the owner of a file or directory, or have root access, to change its permissions.

Here we do not know if brian owns `db_data`. Note:

Permission 7 full

6 read and write

5 read and execute 4 read only

3 write and execute 2 write only

1 execute only

0 none

0 --- no permission 1 --x execute

2 -w- write

3 -wx write and execute 4 r-- read

5 r-x read and execute 6 rw- read and write

7 rwx read, write and execute

Solaris: Solaris Advanced User's Guide

NEW QUESTION 18

Which files must be edited in order to set up logging of all failed login attempts?

- A. `/etc/default/login`, `/var/adm/loginlog`, `/etc/syslog.conf`
- B. `/etc/default/login`, `/var/adm/authlog`, `/etc/syslog.conf`
- C. `/var/adm/loginlog`, `/var/adm/authlog`, `/etc/syslog.conf`
- D. `/etc/default/login`, `/var/adm/authlog`, `/var/adm/loginlog`

Answer: B

Explanation:

This procedure captures in a `syslog` file all failed login attempts.

1. Set up the `/etc/default/login` file with the desired values for `SYSLOG` and `SYSLOG_FAILED_LOGINS`. Edit the `/etc/default/login` file to change the entry. Make sure that `SYSLOG=YES` is uncommented.
2. Create a file with the correct permissions to hold the logging information. Create the `authlog` file in the `/var/adm` directory.
3. Edit the `syslog.conf` file to log failed password attempts. Send the failures to the `authlog` file.

NEW QUESTION 21

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 online `http://pkg.oracle.com/solaris/release`

The update is available on the updated publisher: `PUBLISHERTYPESTATUSURI`

Solaris origin online `http://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 `svccfg - 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.com`
- 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 25

You need to connect two nonglobal zones using a private virtual network. Identify the network resources required in the global zone to accomplish this.

- A. an etherstub and two virtual network interfaces
- B. a virtual bridge
- C. two virtual network interfaces.
- D. two etherstubs

Answer: A

NEW QUESTION 29

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 34

Which two options are accurate regarding the non-global zone console?

- A. Access the non-global zone console by using the zlogin -c command.
- B. Access the non-global zone console by using the zlogin -1 command.
- C. Disconnect from the non-global zone console by using the ~. keys.
- D. Disconnect from the non-global zone console by using the #. keys.

Answer: AC

Explanation:

A: How to Log In to the Zone Console Use the zlogin command with the -C option and the name of the zone, for example, my-zone.

```
global# zlogin -C my-zone
```

C: To disconnect from a non-global zone, use one of the following methods.

* To exit the zone non-virtual console: zonename# exit

* To disconnect from a zone virtual console, use the tilde (~) character and a period: zonename# ~.

NEW QUESTION 39

A user on the system has started a process, but it needs to be terminated. The process ID was determined as follows:

```
pgrep userprogram l5317
```

The user attempted to terminate the program as follows: pkill 15317

This command runs without an error message, and the process continues to run. What is the issue?

- A. You need to run the pkill command with the process name.
- B. You need to switch to super user to kill the process.
- C. You need to run the ps command to get more information.
- D. You need to run the prstat command to get more information.

Answer: B

Explanation:

You can use the pgrep and pkill commands to identify and stop command processes that you no longer want to run. These commands are useful when you

mistakenly start a process that takes a long time to run.

To terminate a process:

Type `pgrep` to find out the PID(s) for the process(es). Type `pkill` followed by the PID(s).

You can kill any process that you own. Superuser can kill any process in the system except for those processes with process IDs of 0, 1, 2, 3, and 4. Killing these processes most likely will crash the system.

NEW QUESTION 41

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 so1.repo.iso-b > so1.full.iso` Mount 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 -G http://pkg.oracle.com/solaris/release/ \-g http://serverA.example.com/ solaris`
- B. `cat so1.repo.iso-a so1.repo.iso-b > /export/IPS/repo` Set the `pkg/inst_root` property to `true` and the `pkg/readonly` property to `/export/IPS` Set the preferred publisher by using `pkg set-publisher -G http://serverA.example.com/ \-g http://pkg.oracle.com/solaris/release/solaris`
- C. `cat so1.repo.iso-a so1.repo.iso-b > so1.full.iso` Mount the ISO image and use the `rsync` command to extract the contents of the ISO file to `/export/IPS/repo` 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 solaris \-g http://pkg.oracle.com/`
- D. `cat so1.repo, iso-a so1.repo.iso-b > /export/IPS/repo.iso` Mount the ISO image and copy the `repo` directory from the ISO image to `/export/IPS/repo` set the `pkg/inst_root` property and the `pkg/readonly` property to `/export/IPS/repo` set 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 44

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 47

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 servic
- 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 servic
- 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 51

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/SMP-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 52

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

It appears to work correctly most of the time, but at other times does not respond. You suspect that the assigned address may not be unique within the network. What command will be useful to confirm this?

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

Answer: E

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 57

User jack makes use of the bash shell; his home directory is /export/home/jack.

What is the correct setting of umask, and where should it be set, to allow jack to create a shell script using the vi editor, that is executable by default?

- A. It is not possible to make a script executable without using the chmod command.
- B. umask value of 0002 set in /etc/profile
- C. umask value of 0002 set in /export/home/jack/.bashrc
- D. umask value of 0722 set in /etc/profile
- E. umask value of 0722 set in /export/home/jack/.bashrc

Answer: B

Explanation:

The user file-creation mode mask (umask) is used to determine the file permission for newly created files. It can be used to control the default file permission for new files. It is a four-digit octal number.

You can setup umask in /etc/bashrc or /etc/profile file for all users. By default most Unix distro set it to 0022 (022) or 0002 (002).

1. The default umask 002 used for normal user. With this mask default directory permissions are 775 and default file permissions are 664.
2. The default umask for the root user is 022 result into default directory permissions are 755 and default file permissions are 644.
3. For directories, the base permissions are (rwxrwxrwx) 0777 and for files they are 0666 (rw-rw-rw).

In short,

1. A umask of 022 allows only you to write data, but anyone can read data.
2. A umask of 077 is good for a completely private system. No other user can read or write your data if umask is set to 077.
3. A umask of 002 is good when you share data with other users in the same group. Members of your group can create and modify data files; those outside your group can read data file, but cannot modify it. Set your umask to 007 to completely exclude users who are not group members.

NEW QUESTION 59

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

```
CHECK all IGNORE dirmtime
/etc/security
/etc/notices IGNORE contents
/export/home
IGNORE mtime size contents
/var CHECK
```

You are using BART to detect inappropriate changes to the file system. Identify the two correct statements describing the attributes recorded.

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

Answer: DF

Explanation:

D: According to line /etc/security F: According to line /export/home

Not E: According to line IGNORE dirmtime

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

```
CHECK all IGNORE dirmtime
```

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

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

NEW QUESTION 63

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/sharstab 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 64

How are operating system updates distributed in the Oracle Solaris 11 environment?

- A. Updates are only available to customers with an active support contract
- B. The updates are distributed through the My Oracle Support web portal and installed in a central location
- C. All software packages are then updated manually from the command line using the smpatch command.
- D. Patches are downloaded from <http://support.oracle.com> either automatically or manually
- E. All software packages are then updated manually from the command line using the smpatch or patchadd commands.
- F. Software updates are published as packages to a repository
- G. All software packages are then updated manually from the command line using the pkg command.
- H. Software updates, published as packages to an OS image
- I. All software packages are then updated manually from the command line using the pkg command.

Answer: C

Explanation:

* Updating all of the packages on your installed system – To update all of the packages on your system that have available updates, use the pkg update command, as follows:

pkg update

Running this command updates packages that you might not otherwise consider updating, for example, kernel components and other low-level system packages.

* Adding or updating individual packages – To add individual software packages, use the pkg install command. Any dependent packages are also updated at the same time.

* install package updates that deliver fixes – A pkg update operation might include bug fixes, so the operation is similar to applying a specific patch or patches in previous Oracle Solaris releases.

Note: The IPS interfaces first check for updates for currently installed packages before retrieving them via the network. By default, interfaces check repository catalogs in the following locations:

* The default installation repository at pkg.oracle.com/solaris/release.

* The support repository in My Oracle Support. This repository is restricted to users with Oracle Solaris 11 Express support contracts, and it contains packages with the latest bug fixes. For this reason, a support contract must be purchased for production deployments.

NEW QUESTION 67

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 71

You have edited /etc/profile to include the lines: dennis_says=hello

export dennis_says

You have also edited /etc/skel/local.profile to include the line: dennis_says=world

You now create a new user account brian, and specify use of the bash shell. When brian logs in and enters

Echo \$dennis_says

What will he see, and why?

- A. world, because the local.profile entry will be executed last
- B. hello, because the global /etc/profile entry overrides the local.profile entry
- C. hello, because the local.profile entry is not automatically sourced on login
- D. hello, because the value specified in local.profile was not exported
- E. nothing, because the variable was not exported in local.profile

Answer: A

Explanation:

The \$HOME/.profile file is an initialization file that is executed after the /etc/profile when logging in to the Bourne or Korn shell. The file contains user preferences for variable settings. If the ENV variable is set to .kshrc, the .kshrc file executes every time a new shell

begins execution. The \$HOME/.profile is copied from the /etc/skel/local.profile file by the Administration Tool when creating a new account.

Note: /etc/skel/local.profile

Per-system configuration file for sh/ksh/ksh93/bash login sessions, installed for new users

NEW QUESTION 74

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

- A. You can create only one manifest per install servic
- B. If you need more than one manifest create multiple install services.
- C. If two client machines have different architectures and need to be installed with the same version of the Oracle Solaris 11 OS, then create two AI manifests and a single install service.
- D. You need a separate install service for each different client architecture that you plan to install, and for each different version of the Oracle Solaris 11 OS that you plan to install on client systems.
- E. If two client machines have different architectures and need to be installed with different versions of the Oracle Solaris 11 OS, then create two AI manifests and two install services.
- F. The install server needs to be able to access an Oracle Solaris Image Packaging System (IPS) software package repository; the clients do not.
- G. The install server can be either an x86 machine or a SPARC machine.

Answer: BEF

Explanation:

B (not A, not D, Not C): If two client machines need to be installed with the same version of the Oracle Solaris 11 OS but need to be installed differently in other ways, then create two AI manifests for the AI install service. The different AI manifests can specify different packages to install or a different

slice as the install target, for example.

Note: An AI manifest provides installation instructions.

The AI manifest specifies one or more IPS package repositories where the client retrieves the packages needed to complete the installation. The AI manifest also includes the names of additional packages to install and information such as target installation device and partition information.

F: The install server can be either an x86 machine or a SPARC machine.

NEW QUESTION 77

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 81

Which two options are valid methods of installing a solaris10 branded zone on a system running Oracle Solaris 11?

- A. Use the V2V process to migrate an existing Solaris 8 or 9 non-global zone from a Solaris 10 system to a solaris10 branded zone.
- B. Use the V2V process to migrate an existing Solaris 10 non-global whole root zone from a Solaris 10 system to a solaris10 branded whole root zone.
- C. Install a solaris10 branded zone directly from the Oracle Solaris 10 media.
- D. Migrate an existing 64-bit Solaris 10 system to a solaris10 branded non-global zone using the P2V process.
- E. Use the V2V process to migrate an existing Solaris 10 non-global sparse root zone from a Solaris 10 system to a solaris10 branded sparse root zone.

Answer: BC

Explanation:

B: How to Migrate an Existing native Non-Global Zone

Use the V2V process to migrate an existing zone on your Solaris 10 system to a solaris10 brand zone on a system running the Oracle Solaris 11 release.

C: How to Install the solaris10 Branded Zone

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

NEW QUESTION 86

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 89

User jack, whose account is configured to use the korn shell, logs in and examines the value of his PATH environment variable:

```
jack@solaris: echo $PATH
/usr/gnu/bin:/usr/bin:/usr/sbin:/sbin
```

There is a shell script in jack's home directory called useradd:

```
-r-xr-xr-x 2 jack other 1239 2012-01-05 11:42 useradd
```

While in his home directory, jack attempts to run the script:

```
jack@solaris: useradd
```

What will happen, and why?

- A. He will get a "file not found" error, because the current directory is not in his search path.
- B. He will get a "file not found" error, because his home directory is not in his search path.
- C. The useradd script will execute, because jack is in the same directory that the script is located in.
- D. The command /user/sbin/useradd will execute, because it is the last match in the search path.
- E. The command /user/sbin/useradd will execute, because it is the first match in the search path.

Answer: D

NEW QUESTION 91

user1, while in his home directory, is attempting to run the following command in his home directory: cp bigfile verybig

The system displays the following error:

```
cp: cannot create verybig: Disc quota exceeded
```

Your initial troubleshooting shows that the df -h command indicates he is at 100% capacity. What command would you use to increase the disk space available to the user?

- A. zfs get quota rpool/export/home/user1
- B. zfs userused@user1
- C. zfs quota=none /rpool/export/home/user1
- D. df -h | grep user1
- E. zfs set quota=none /rpool/export/home/user1

Answer: E

Explanation:

ZFS quotas can be set and displayed by using the zfs set and zfs get commands. We can remove the quota restriction by setting to quota to none.

NEW QUESTION 96

Which two are true about accounts, groups, and roles in the Solaris user database?

- A. All Solaris user accounts must have a unique UID number.
- B. A Solaris account name may be any alphanumeric string, and can have a maximum length of 8 characters.
- C. Account UID numbers 0-09 are system-reserved.
- D. The GID for an account determines the default group ownership of new files created by that account.
- E. The groups that an account is a member of are determined by the entries in the/etc/group file.

Answer: AB

Explanation:

A: Solaris uses a UID (User ID) to identify each user account. The UID is a unique number assigned to each user. It is usually assigned by the operating system when the account is created.

B: In Solaris the account name can include any alphanumeric string (and . _ -). The maximum length is 8 characters.

NEW QUESTION 100

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 104

A user brian is configured to use the bash shell. His home directory is /export/home/brian, and contains a .profile and a .bashrc file.

In the .profile, there are these lines: genius =ritchie

export genius

In the .bashrc us this line: genius=kernighan

In /etc/profile are these lines: genius=thompson

export genius

When brian logs in and asks for the value of genius, what will he find, and why?

- A. genius will be ritchie, because that was the value exported in .profile.
- B. genius will be kernighan, because .bashrc executes after .profile.
- C. genius will be ritchie because variable settings in .profile take precedence over variable settings in .bashrc.
- D. genius will be ritchie because .profile executes after .bashrc.
- E. genius will be thompson because /etc/profile system settings always override local settings.

Answer: C

NEW QUESTION 105

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

Review the information taken from your server:

```
rpool@BE1
rpool/ROOT@BE1
rpool/ROOT/solaris@BE1
rpool/ROOT/dump@BE1
rpool/ROOT/export@BE1
rpool/ROOT/export/home@BE1
rpool/ROOT/swap@BE1
```

Which option describes the command used to create these snapshots of the root file system?

- A) `zfs snapshot -r rpool@BE1`
- B) `beadm create -n BE1`
- C) `zfs snapshot -r BE1 rpool`
- D) `zfs snapshot rpool BE1`
- E) `zfs snapshot rpool@BE1 rpool/ROOT@BE1 rpool/ROOT/solaris@BE1 \`
`rpool/ROOT/dump@BE1 rpool/ROOT/export@BE1 \`
`rpool/ROOT/export/home@BE1 rpool/ROOT/swap@BE1`

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

Answer: A

Explanation:

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

`-r`

Recursively create snapshots of all descendent datasets. Snapshots are taken atomically, so that all recursive snapshots correspond to the same moment in time.

NEW QUESTION 109

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 113

In Oracle Solaris 11, where is the Oracle default repository located?

- A. `/var/spool/pkg`
- B. `http://localhost/solaris`
- C. `http://pkg.oracle.com/solaris/release`
- D. `http://www.oracle.com/Solaris/download`
- E. `/cdrom/cdrom0`

Answer: C

Explanation:

REPOSITORY DESCRIPTION

* `http://pkg.oracle.com/solaris/release`

The default repository for new Oracle Solaris 11 users. This repository receives updates for each new release of Oracle Solaris. Significant bug fixes, security updates, and new software may be provided at any time for users to install at Oracle's discretion.

* `https://pkg.oracle.com/solaris/support`

Provides bug fixes and updates. Accessible with a current support contract from Oracle.

* `https://pkg.oracle.com/solaris/dev` Provides the latest development updates. Accessible to users enrolled in the Oracle Solaris 11 Platinum Customer Program and approved Oracle Partners.

NEW QUESTION 115

Identify the correctly matching pair of equivalent functionality of JumpStart and Automated installer (AI).

- A. JumpStart: `begin` script AI: package repository
- B. JumpStart: `setup_server` AI: `installadm create-service`
- C. JumpStart: `add_install_client` AI: SMF system configuration profile files
- D. JumpStart: finish scripts and sysidsfg files AI: manifest files

Answer: B

Explanation:

JumpStart: Use the `setup_install_server(1M)` command. AI: Use the `installadm create-service` command.

NEW QUESTION 119

What is the output of the following command, if executed using the default shell for the root role account of a standard Live CD Install of Oracle Solaris 11?
echo '\$SHELL'

- A. /usr/bin/bash
- B. /usr/bin/ksh
- C. \$SHELL
- D. the PID for the current shell

Answer: C

Explanation:

Single quotes are most strict. They prevent even variable expansion. Double quotes prevent wildcard expansion but allow variable expansion. For example:
#!/bin/sh echo \$SHELL
echo "\$SHELL"
echo '\$SHELL' This will print:
/usr/bin/bash
/usr/bin/bash
\$SHELL

NEW QUESTION 120

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 125

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 128

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

You have these requirements:

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

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

- A. Svccfg - s filesystem/shadowd:defaultsetprop config_params/shadow_threads=integer: 16endsvcadm refresh filesystem/shadowd: default
- B. Specify the -b <blocksize> option with the zfs create command and increase the value of <blocksize>
- C. Use the -o -volblocksize=<blocksize>option with the zfs create command and increase the value of the default <blocksize>.
- D. Svccfg -s filesystem/zfs: defaultsetprop config_params/shadow_threads = integer: 16endsvcadm refresh filesystem/zfs:default

Answer: A

Explanation:

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

The shadowd service is managed by the service management facility, smf(5).

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). The service's status can be queried using the svcs(1) command.

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

config_params/shadow_threads

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

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

file:///path nfs://host:path

NEW QUESTION 130

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 131

Your system is assigned an IP address object 192.168.0.222/24. However, the net mask — expressed as four octets — is required. Which is the correct netmask?

- A. 255.0.0.0
- B. 255.255.0.0
- C. 255.255.255.0
- D. 255.255.255.24
- E. 255.255.255.255

Answer: C

Explanation:

A 24-bit network mask is expressed as 255.255.255.0.

NEW QUESTION 135

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 137

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 138

User1 is attempting to run the following command: cp bigfile verybig

The system displays the following error:

cp: cannot create verybig: Disc quota exceeded

Your initial troubleshooting shows that the df -h command indicates the account is at 100% capacity. What command would you use to determine how much disk space the user has available?

- A. zfs get quota rpool/export/home/user1
- B. zfs userused@user1
- C. zfs quota=1M /rpool/export/home/user1
- D. df -h | grep user1

Answer: A

Explanation:

ZFS quotas can be set and displayed by using the zfs set and zfs get commands. In the following example, a quota of 10 Gbytes is set on tank/home/bonwick.

```
# zfs set quota=10G tank/home/bonwick
# zfs get quota tank/home/bonwick NAME PROPERTY VALUE SOURCE
tank/home/bonwick quota 10.0G local
```

NEW QUESTION 139

Which two statements are true concerning the network stack on Oracle Solaris 11?

- A. Hardware network interfaces and datalinks have a one-to-one relationship.
- B. IP addresses are assigned to datalinks.
- C. A single IP interface can have either an IPv4 address or an IPv6 address but not both.
- D. A single IP interface can have both an IPv4 address and an IPv6 address.
- E. A single datalink can have only one IP interface.

Answer: AD

NEW QUESTION 143

You are troubleshooting a newly installed desktop Oracle Solaris 11 system with a single network interface. From this system, you can connect to other systems within the company

intranet, but cannot access any external services (such as websites and email), even when using IP addresses.

Examining the routing table confirms that the default route to 192.168.1.1 is missing. DHCP is not used at this site. Which two commands will temporarily and permanently configure the default route?

- A. ipadm set-gateway 192.168.1.1
- B. route add default 192.168.1.1
- C. ipadm set-default 192.168.1.1
- D. dladm route-add -d 192.168.1.1
- E. echo 192.168.1.1 >/etc/gateway
- F. echo 192.168.1.1 >/etc/defaultrouter

Answer: BF

Explanation:

B: Setting the default route on Solaris is easy. If you are trying to just set the route temporarily you can use the route command:

Route add default <ipaddress> Example:

```
Route add default 192.168.1.1
```

Note: Route command manipulates the kernel routing tables. Routing is the process of forwarding a packet from one computer to another. It is based on the IP address in the IP packet header and netmask.

F: If you want the route to be persisted when you reboot the system, you will need to set the route in the /etc/defaultrouter file.

/etc/defaultrouter Example:

```
Echo 192.168.1.1 > /etc/defaultrouter
```

NEW QUESTION 145

The global zone has 8 CPUs. YOU suspect that one of your non global /ones, dbzone, is consuming all of the CPU resources.

Which command would you use to view the CPU utilization for all of the zones to confirm this?

- A. Run from the global zone:prstat -Z
- B. Run from each zone:zlogin <zonename> mpstat
- C. Run from the global zone:zonestar -r summary
- D. Run from the global zone:rctladm -l

E. Run from the global zone:prctl -i

Answer: A

Explanation:

If you're logged on to the system, you can run prstat -Z to generate a summary of cpu/memory utilization by zone.

NEW QUESTION 148

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

- A. While configuring the zone, add this entry:add rct1set name = capped.cpu-sharesadd value (priv = privileged, limit = 20, action = none)endexit
- B. While configuring the zone, add this entry: add rct1set name= zone.cpu-sharesadd value (priv=privileged, limit=20, action=none)endexitfrom command line, enter: # dispadmin- d FSS
- C. From the command line enter: #prctl -n zone.cpu-shares - r - v 20 - i zone testzone
- D. From the command line, enter:#prctl - n zone.cpu-shares - v 80 - r - i zone global

Answer: C

Explanation:

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

How to Change the zone.cpu-shares Value in a Zone Dynamically This procedure can be used in the global zone or in a non-global zone.

For more information about roles, see Configuring and Using RBAC (Task Map) in System Administration Guide: Security Services.

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

idtype is either the zonename or the zoneid. value is the new value.

Note: project.cpu-shares

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

NEW QUESTION 150

The following image properties are displayed on your system:

PROPERTY	VALUE
be-policy	always-new
ca-path	/etc/openssl/certs
check-certificate-revocation	False
flush-content-cache-on-success	True
mirror-discovery	False
preferred-authority	
publisher-search-order	['solaris']
send-uuid	True
signature-policy	verify
signature-required-name	[]
trust-anchor-directory	etc/certs/CA
use-system-repo	False

Which two options describe the boot environment policy property that is currently set for this image?

- A. All package operations are performed in a new BE set as active on the next boot.
- B. Do not create a new B
- C. The install, update, uninstall, or revert operation is not performed if a new BE is required.
- D. If a BE is created, do not set it as the active BE on the next boot
- E. A reboot is required for all package operations
- F. A reboot is not required after a package operation.
- G. For package operations that require a reboot, this policy creates a new BE set as active on the next boot.

Answer: DF

Explanation:

Image properties described below.

* be-policy

Specifies when a boot environment is created during packaging operations. The following values are allowed:

/ default

Apply the default BE creation policy: create-backup.

/ always-new (D, F)

Require a reboot for all package operations (D) by performing them in a new BE set as active on the next boot (F). A backup BE is not created unless explicitly requested.

This policy is the safest, but is more strict than most sites need since no packages can be added without a reboot.

NEW QUESTION 155

Oracle Solaris 11 limits access to the system with usernames and passwords.

The usernames are held in , and the passwords are held in . Select the correct pair.

- A. /etc/security/policy.conf /etc/passwd
- B. /etc/passwd /etc/shadow
- C. /etc/security /etc/passwd
- D. /etc/shadow /etc/passwd

Answer: B

Explanation:

The /etc/passwd file contains basic user attributes. This is an ASCII file that contains an entry for each user. Each entry defines the basic attributes applied to a user.

/etc/shadow file stores actual password in encrypted format for user's account with additional properties related to user password i.e. it stores secure user account information. All fields are separated by a colon (:) symbol. It contains one entry per line for each user listed in /etc/passwd file.

NEW QUESTION 159

Examine this command and its output:

```
$ zfs list -r -t all tank
Name USED AVAIL REFER MOUNTPOINT
tank 2.41G 2.43G 32K /tank
tank/database 2.41G 2.43G 2.41G /tank/database tank/[email protected] 20K - 2.00G -
```

Next you execute:

```
# zfs destroy tank/database
```

Which statement is true about the result of executing this command?

- A. It destroys the tank/database dataset.
- B. It destroys tank/database and all descendant datasets.
- C. It fails because the tank/ snapshot depends on the tank/database dataset.
- D. It fails because the tank/ clone depends on the tank/database dataset.
- E. It fails because the tank/database data set is not empty.

Answer: C

NEW QUESTION 164

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 166

You create a flash archive of the Solaris 10 global zone on the server named sysA. The archive name is s10-system.flar, and it is stored on a remote server named backup_server.

On sysA, you create a Solaris 10 branded zone named s10-zone.

You want to use the flash archive, located on /net/backup_servers/10-system.flar, to install the Operating system in the s10-zone zone.

Which command do you choose to install the s10-system.flar archive in the Solaris 10 branded zone (s10-zone)?

- A. zoneadm -z s10 -zone install -a /net/backup_server/s10-system.flar -u
- B. zonecfg -z s10 -zone install -a /net/backup_server/s10-system.flar -u
- C. zoneadm -z s10 -zone clone -s /net/backup_server/s10-system.flar
- D. zonecfg -a s10-zone create -t SUNWsolaris10\</net/backup_server/s10-system.flar
- E. zonecfg -z s10-zone install -f /net/backup/backup_server/s10-system.flar

Answer: A

Explanation:

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

How to Install the solaris10 Branded Zone

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

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

NEW QUESTION 169

User jack logs in to host solaris 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
K3eropfqsmPPY6uilyh9vpiFX2r1LHcbx139+vG6HOtVvuS8+lfMDtawm3WQvRuOopz3vVy
5GtMwtaOgehsXoT930Ryev1bH5myPtWKlipITsOd2sX9k3tvjrmme4KCGGss=
```

NEW QUESTION 174

Which two statements are true when updating Solaris 11 from one Support Respository Update (SRU) to another SRU by using the pkg update command?

- A. By default, the pkg update command automatically creates a backup Boot Environment whenever the kernel is affected by the update.
- B. By default, the pkg update command automatically creates a new Boot Environment whenever the kernel is affected by the update.
- C. The pkg update command can only be used to update to a newer SRU.
- D. The pkg update command can be used to update to a newer or older SRU.
- E. By default, the pkg update command always updates Solaris 11 to the first SRU that was released after the Current SRU.
- F. The pkg update command can only be performed while running in the single-user milestone.

Answer: BC

NEW QUESTION 175

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. A virtual router
- C. A virtual switch
- D. A virtual bridge.
- E. A virtual network interface
- F. Nothing because a virtual switch is automatically created then the virtual network interfaces are created.

Answer: A

Explanation:

Etherstubs are pseudo ethernet NICs which are managed by the system administrator. You can create VNICs over etherstubs instead of over physical links. VNICs over an etherstub become independent of the physical NICs in the system. With etherstubs, you can construct a private virtual network that is isolated both from the other virtual networks in the system and from the external network. For example, you want to create a network environment whose access is limited only to your company developers than to the network at large. Etherstubs can be used to create such an environment.

Note: Oracle Solaris 11 introduces a new and powerful network stack architecture which includes:

* Networking virtualization with virtual network interface cards (VNICs) and virtual switching (etherstubs)

* Tight integration with zones

* Network resource management - efficient and easy to manage integrated quality of service (QoS) to enforce bandwidth limit on VNICs and traffic flows

NEW QUESTION 180

Given:

file1 and file2 are text files. dir1 and dir2 are directories.

Which two commands will be successful?

- A. cp dir1 dir1
- B. cp dir1 file1
- C. cp file? dir1
- D. cp fil
- E. dir1
- F. cp file% dir2
- G. cp file1 file2 dir1

Answer: CF

Explanation:

C: Here the wildcard character ? is used (Matches any single character). file1 and file2 will be copied into dir1

F: the two files file1 and file2 are copied into directory dir1. Note: cp - copy files and directories

Copy SOURCE to DEST, or multiple SOURCE(s) to DIRECTORY.

Cp has three principal modes of operation. These modes are inferred from the type and count of arguments presented to the program upon invocation.

* When the program has two arguments of path names to files, the program copies the contents of the first file to the second file, creating the second file if necessary.

* When the program has one or more arguments of path names of files and following those an argument of a path to a directory, then the program copies each source file to the destination directory, creating any files not already existing.

* When the program's arguments are the path names to two directories, cp copies all files in the source directory to the destination directory, creating any files or directories needed. This mode of operation requires an additional option flag, typically r, to indicate the recursive copying of directories. If the destination directory already exists, the source is copied into the destination, while a new directory is created if the destination does not exist.

NEW QUESTION 182

Identify three options that describe the new Oracle Solaris 11 zone features.

- A. There are boot environments for zones.
- B. Administrators can delegate common administration tasks by using RBAC.
- C. Oracle Solaris 11 supports Solaris 8, 9, and 10 branded zones.
- D. You can migrate a physical Solaris 10 system and its non-global zones to a solaris10 branded zone running on an Oracle Solaris 11 system.
- E. It is possible to change the host ID of a zone.

Answer: ABD

Explanation:

A: The beadm utility includes support for creating and administering non-global zone boot environments.

Note: A boot environment is a bootable instance of the Oracle Solaris operating system image plus any other application software packages installed into that image. System administrators can maintain multiple boot environments on their systems, and each boot environment can have different software versions installed.

B: Role-based access control (RBAC) is a security feature for controlling user access to tasks that would normally be restricted to the root role. By applying security attributes to processes and to users, RBAC can divide up superuser capabilities among several administrators.

NEW QUESTION 185

You notice that the /var/.dm/messages file has become very large. Typically, this is managed by a crontab entry. Which entry should be in the root's crontab file?

- A. 10 3 * * * /usr/adm/messages
- B. 10 3 * * * /usr/sbin/logadm
- C. 10 3 * * * /usr/sbin/syslogrotate
- D. 10 3 * * * /usr/sbin/logrotate
- E. 10 3 * * * /usr/sbin/messages

Answer: B

Explanation:

This example shows how to display the default root crontab file.

```
$ suPassword:
```

```
# crontab -l
```

```
#ident "@(#)root 1.19 98/07/06 SMI" /* SVr4.0 1.1.3.1 */
```

```
#
```

```
# The root crontab should be used to perform accounting data collection.
```

```
#
```

```
#
```

```
10 3 * * * /usr/sbin/logadm
```

```
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
```

```
30 3 * * * [ -x /usr/lib/gss/gsscred_clean ] && /usr/lib/gss/gsscred_clean
```

```
#10 3 * * * /usr/lib/krb5/kprop_script slave_kdcs
```

NEW QUESTION 186

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 188

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 193

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 198

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 200

The /usr/bin/p7zip file that is part of the p7zip package has been overwritten. This server is critical to production and cannot be rebooted. Identify the command that would restore the file without requiring a reboot.

- A. pkg verify p7zip
- B. pkg fix p7sip
- C. pkg rebuild-index p7zip
- D. pkg revert p7zip
- E. pkg uninstll p7zip
- F. pkg install p7zip
- G. pkg install --no-backup-be p7sip
- H. pkg refresh p7zip

Answer: D

Explanation:

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

NEW QUESTION 205

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 208

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 209

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

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

Answer: CE

NEW QUESTION 210

View the following information for a software package:

```
Name: compress/gzip
Summary: GNU zip (gzip)
Description: The GNU Zip (gzip) compression utility
Category: Applications/System utilities
State: Installed
Publisher: Solaris
Version: 1.3.5
Build Release: 5.11
Branch: 0.175.0.0.0.2-537
Packaging Date: October 19, 2011 09:12:46 AM
Size: 215.32 kB
FMRI:
pkg://solaris/compress/gzip@1.3.5,5.11-0.175.0.0.0.2.537:20111019T091246z
```

Which command would you use to display this information for a software package that is not currently installed on your system?

- A. pkg list gzip
- B. pkg info -r gzip
- C. pkg search -1 gzip

- D. pkg verify -v gzip
- E. pkg contents gzip

Answer: B

Explanation:

By default, the pkg info command only lists information about installed packages on the system; however, we can use a similar command to look up information about uninstalled packages, as shown in here:

Example:

Listing Information About an Uninstalled Package

```
# pkg info -r php-52 Name: web/php-52 Summary: PHP Server 5.2
```

```
Description: PHP Server 5.2 Category: Development/PHP State: Not Installed Publisher: solaris
```

```
Version: 5.2.17
```

```
Build Release: 5.11
```

```
Branch: 0.175.0.0.0.1.530
```

```
Packaging Date: Wed Oct 12 14:01:41 2011
```

```
Size: 44.47 MB
```

```
FMRI: pkg://solaris/web/php-52@5.2.17, 5.11-0.175.0.0.0.1.530:20111012T140141Z
```

Note: pkg info command displays information about packages in a human-readable form. Multiple FMRI patterns may be specified; with no patterns, display information on all

installed packages in the image.

With -l, use the data available from locally installed packages. This is the default.

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 215

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 220

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 222

You wish to edit your crontab file that is located in /var/spool/cron/crontab. What command must you enter to edit this file?

- A. crontab -e
- B. crontab -e /var/spool/cron/crontab
- C. crontab -r
- D. crontab -e /etc/default/cron

Answer: A

Explanation:

The main tool for setting up cron jobs is the crontab command, though this is not available on every Unix variant. Typically under Solaris or Linux one would create a new crontab or edit an existing one, using the command;

crontab -e

Use the ls -l command to verify the contents of the /var/spool/cron/crontabs file.

NEW QUESTION 226

Select the packet type that identifies members of the group and sends information to all the network interfaces in that group.

- A. Unicast
- B. Multicast
- C. Broadcast
- D. Bayesian
- E. Quality of Service Priority

Answer: B

Explanation:

IPv6 defines three address types: unicast

Identifies an interface of an individual node.

multicast

Identifies a group of interfaces, usually on different nodes. Packets that are sent to the multicast address go to all members of the multicast group.

anycast

Identifies a group of interfaces, usually on different nodes. Packets that are sent to the anycast address go to the anycast group member node that is physically closest to the sender.

NEW QUESTION 231

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 233

You are installing the Oracle Solaris 11 Operating System by using the Text Installer. Which two options describe the features associated with the Text Installer?

- A. It can be used to install only SPARC systems.
- B. It installs gnome as the default user environment on a system capable of displaying a graphical environment.
- C. You can choose whether root is a role or user account.
- D. You can do both automatic and manual configuration of the network.
- E. You can select how to configure the remaining network interfaces.

Answer: CD

NEW QUESTION 234

View the Exhibit.

```

ascii name = <ATA-VBOX HARDDISK-1.0-16.00GB>
bytes/sector = 512
sectors = 33554431
accessible sectors = 33554398
Part      Tag      Flag      First Sector      Size      Last Sector
 0         usr       wm          256             15.99GB   33538014
 1 unassigned wm           0                0          0
 2 unassigned wm           0                0          0
 3 unassigned wm           0                0          0
 4 unassigned wm           0                0          0
 5 unassigned wm           0                0          0
 6 unassigned wm           0                0          0
 8 reserved wm          33538015         8.00MB    33554398
format>

```

Which is true regarding the disk drive?

- A. This disk configuration could be used as a ZFS root disk.
- B. This disk contains an SMI disk label.
- C. Slice 7 represents the entire disk and cannot be used as a slice for a file system
- D. The disk contains an EFI disk label.

Answer: A

Explanation:

Installing a ZFS Root Pool

The installer searches for a disk based on a recommended size of approximately 13 GB.

NEW QUESTION 237

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

NAME (PUBLISHER)VERSIONINFO

Compress/zip3.1.2-0.175.0.0.0.0.537if-

NAMEVERSIONDATECOMMENT

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

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

- A. This package cannot be removed.
- B. This package can be updated to a new version when the new version of the package becomes available.
- C. This package cannot be updated.
- D. This package can be updated to version 3.1.3 but not 3.2.
- E. This package cannot be downgraded to version 3.1.1.

Answer: B

Explanation:

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

Note: The "i" in the I column indicates that these packages are installed in this image. Adding and Updating Oracle Solaris 11 Software Packages, Showing Package Install State Information

NEW QUESTION 239

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 241

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 242

Which five statements describe options available for installing the Oracle Solaris 11 operating system using the installation media?

- A. You can perform a text or LiveCD installation locally or over the network.
- B. The text Installer does not install the GNOME desktop.
- C. The GNOME desktop package must be added after you have installed the operating system.
- D. The LiveCD Installation cannot be used to install multiple instances of Oracle Solaris.
- E. The LiveCD installer cannot be used if you need to preserve a specific Solaris Volume Table of Contents (VTOC) slice in your current operating system.
- F. The LiveCD Installer is for x86 platforms only.
- G. The GUI installer cannot be used to upgrade your operating system from Solaris 10.
- H. If you are installing Oracle Solaris 11 on an x86-based system that will have more than one operating system installed in it, you cannot partition your disk during the installation process.
- I. The LiveCD installer can be used for SPARC or x86 platforms.

Answer: ABDFH

Explanation:

A: If the network is setup to perform automated installations, you can perform a text installation over the network by setting up an install service on the network and selecting a text installation when the client system boots.

B: After a fresh install of Solaris 11 express, only the console mode is activated. To add Gnome, simply do :

```
$ sudo pkg install slim_install
```

This will install additional packages that are not installed by default. D: The text installer advantages over the GUI installer include:

* In addition to modifying partitions, the text installer enables you to create and modify VTOC slices within the Solaris partition.

F: How do I upgrade my Solaris 10 or lower systems to Solaris 11?

Unfortunately, you CAN'T. There is no direct upgrade installer or other tool that will allow you to upgrade from earlier releases of Solaris to Solaris 11. This is primarily due to the vast changes in the packaging mechanism in Solaris 10.

NEW QUESTION 243

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

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

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

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

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

Answer: A

Explanation:

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

Note: zpool attach [-f] pool device new_device

Attaches new_device to an existing zpool device. The existing device cannot be part of a raidz configuration. If device is not currently part of a mirrored configuration, device automatically transforms into a two-way mirror of device and new_device. If device is part of a two-way mirror, attaching new_device creates a three-way mirror, and so on. In either case, new_device begins to resilver immediately.

NEW QUESTION 245

You are currently working in both your home directory and the system directory /tmp. You are switch back and forth with full path names. Which pair of cd commands will provide you with a shortcut to switch between these two locations?

- A. cd ~ and cd -
- B. cd and cd.
- C. cd ~ and cd
- D. cd * and cd . .

Answer: A

Explanation:

In the Bourne Again, C, Korn, TC, and Z shells, the tilde (~) is used as a shortcut for specifying your home directory.

cd -

It's the command-line equivalent of the back button (takes you to the previous directory you were in).

Note:

To make certain that you are in your home directory, type the cd (change directory) command. This command moves you to your home (default) directory.

NEW QUESTION 250

You are logged in to a Solaris 11 system as user jack. You issue the following sequence of commands:

```
jack@solaris:~$ id
uid=65432 (jack) gid=10(staff) groups=10(staff)
jack@solaris:~$ su
Password:
jack@solaris:~#
```

Identify two correct statements.

- A. You have the effective privilege of the account root.
- B. Your GID is 10.
- C. Your home directory is /root.
- D. You are running the shell specified for the account root.
- E. Your UID is 1.

Answer: AB

Explanation:

Oracle Solaris provides predefined rights profiles. These profiles, listed in the

/etc/security/prof_attr, can be assigned by the root role to any account. The root role is assigned all privileges and all authorizations, so can perform all tasks, just as root can when root is a user.

To perform administrative functions, you open a terminal and switch the user to root. In that terminal, you can then perform all administrative functions.

```
$ su - root
```

Password: Type root password

```
#
```

When you exit the shell, root capabilities are no longer in effect.

NEW QUESTION 252

You display the IP interface information with `ipmpstat -i`.

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

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

Answer: DE

Explanation:

FLAGS

Indicates the status of each underlying interface, which can be one or any combination of the following:

(D) d indicates that the interface is down and therefore unusable.

(E) M indicates that the interface is designated by the system to send and receive IPv6 multicast traffic for the IPMP group.

Note:

i indicates that the INACTIVE flag is set for the interface. Therefore, the interface is not used to send or receive data traffic.

s indicates that the interface is configured to be a standby interface.

m indicates that the interface is designated by the system to send and receive IPv4 multicast traffic for the IPMP group.

b indicates that the interface is designated by the system to receive broadcast traffic for the IPMP group.

h indicates that the interface shares a duplicate physical hardware address with another interface and has been taken offline. The h flag indicates that the interface is unusable.

NEW QUESTION 256

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 261

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 265

You want to display network interface information. Which command should you use?

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

Answer: A

NEW QUESTION 266

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 271

Solaris 11 includes a redesigned software packaging model: the Image Packaging system.

Which three describe advantages of the Image Packaging System over the previous Solaris 10 SVR4 packaging model?

- A. Eliminates patching of the software package
- B. Makes the patching process more efficient with less downtime
- C. Eliminates OS version upgrade
- D. Allows for the installation of the OS without a local DVD or installation server
- E. Allows the use of a repository mirror to speed up package operation
- F. Allows users to publish their own software package in a software repository

Answer: AEF

NEW QUESTION 275

You are setting up a local IPS package repository on your Oracle Solaris11 server: solaris.example.com.

You want to point the existing local IPS publisher to the new local IPS repository located in /repo.

These are the steps that you have followed:

1. Download and rsync the contents of the Oracle Solaris11 repository ISO image to the

/repo directory.

2. Configure the repository server service properties. The svcprop command display, the IPS related properties:

```
pkg/inst_root astring/repo pkg/readonly Boolean true
```

The 1s command displays the contents of the /repo directory:

```
#ls/repo
```

```
Pkg5.repository publisher
```

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

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

The pkg publisher command still displays: PUBLISHERTYPESTATUSURI

```
solarisoriginonlinehttp://pkg.oracle.com/solaris/release/
```

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

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

Answer: C

Explanation:

Set the Publisher Origin To the File Repository URI

To enable client systems to get packages from your local file repository, you need to reset the origin for the solaris publisher. Execute the following command on each client:

Example:

```
# pkg set-publisher -G '*' -M '*' -g /net/host1/export/repoSolaris11/ solaris
```

NEW QUESTION 278

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 281

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 286

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 288

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, lzh, 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 293

Subnets are created by using .

- A. subnet
- B. netmask
- C. unicast
- D. broadcast

Answer: B**Explanation:**

The process of subnetting involves the separation of the network and subnet portion of an address from the host identifier. This is performed by a bitwise AND operation between the IP address and the (sub)network prefix. The result yields the network address or prefix, and the remainder is the host identifier.

The routing prefix of an address is written in a form identical to that of the address itself. This is called the network mask, or netmask, of the address. For example, a specification of the most-significant 18 bits of an IPv4 address, 11111111.11111111.11000000.00000000, is written as 255.255.192.0.

NEW QUESTION 297

You are executing this command in the default shell: sleep 5000 &

The system displays a number. This value is .

- A. the priority of the /usr/bin/sleep process
- B. the process ID of the /usr/bin/sleep process
- C. the process ID of the shell spawned to execute /usr/bin/sleep
- D. the process group ID that includes the /usr/bin/sleep process
- E. the amount of memory allocated to the /usr/bin/sleep process
- F. the current number of instances of the /usr/bin/sleep process

Answer: C**Explanation:**

If a command is terminated by the control operator '&', the shell executes the command asynchronously in a subshell. This is known as executing the command in the background. The shell does not wait for the command to finish, and the return status is 0 (true).

NEW QUESTION 299

Which four can the SMF notification framework be configured to monitor and report?

- A. all service transition states
- B. service dependencies that have stopped or faulted
- C. service configuration modifications
- D. legacy services that have not started
- E. services that have been disabled
- F. service fault management events
- G. processes that have been killed

Answer: AEFG**Explanation:**

Note 1: State Transition Sets are defined as: to<state>

Set of all transitions that have <state> as the final state of the transition.

form-<state>

Set of all transitions that have <state> as the initial state of the transition.

<state>

Set of all transitions that have <state> as the initial state of the transitional. Set of all transitions. (A)

Valid values of state are maintenance, offline (G), disabled (E), online and degraded. An example of a transitions set definition: maintenance, from-online, to-degraded.

F: In this context, events is a comma separated list of SMF state transition sets or a comma separated list of FMA (Fault Management Architecture) event classes. events cannot have a mix of SMF state transition sets and FMA event classes. For convenience, the tags problem- {diagnosed, updated, repaired, resolved} describe the lifecycle of a problem diagnosed by the FMA subsystem - from initial diagnosis to interim updates and finally problem closure.

Note 2:

SMF allows notification by using SNMP or SMTP of state transitions. It publishes Information Events for state transitions which are consumed by notification daemons like snmp-notify(1M) and smtp-notify(1M). SMF state transitions of disabled services do not generate notifications unless the final state for the transition is disabled and there exist notification parameters for that transition. Notification is not be generated for transitions that have the same initial and final state.

NEW QUESTION 300

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 301

You have completed configuring a zone named dbzone on your Solaris 11 server. The configuration is as following:

```
zonename: dbzone
zonepath: /export/dbzone
brand: Solaris
autoboot: false
bootargs:
file-mac-profile:
pool:
limitpriv:
scheduling-class:
ip-type: exclusive
hostid:
fs-allowed:
anet:
    linkname: net0
    lower-link: auto
    allowed-address not specified
    configure-allowed-address: true
    defrouter not specified
    allowed-dhcp-cids not specified
    link-protection: mac-nospoof
    mac-address: random
    mac-prefix not specified
    mac-slot not specified
    vlan-id not specified
    priority not specified
    rxrings not specified
    rxrings not specified
    mtu not specified
    maxlow not specified
    rxfanout not specified
```

The global zone displays the following network information:

ADDROBJ	TYPE	STATE	ADDR
lo0/v4	static	ok	127.0.0.1/8
net0/_b	dhcp	ok	10.0.2.18/24
lo0/v6	static	ok	::1/128
net0/_a	addrconf	ok	fe80::a00:27ff:fe8e:c0d4/10

The zone has never been booted. Which three options correctly describe this zone?

- A. It is a sparse root zone.
- B. It is a whole root zone.
- C. It is an immutable zone.
- D. It is a native zone.
- E. The zone shares the network interface with the host.
- F. The zone uses a virtual network interface.
- G. The hostid is the same as the global zone.
- H. The IP address of the zone is 10.0.2.18.

Answer: CEG

Explanation:

C: Immutable Zones provide read-only file system profiles for solaris non-global zones. Note that ip-type: exclusive:

Starting with OpenSolaris build 37 and Oracle Solaris 10 8/07, a default zone can be configured as an "exclusive-IP zone" which gives it exclusive access to the NIC(s) that the zone has been assigned. Applications in such a zone can communicate directly with the NIC(s) available to the zone.

Note on zones:

After installing Oracle Solaris on a system, but before creating any zones, all processes run in the global zone. After you create a zone, it has processes that are associated with that zone and no other zone. Any process created by a process in a non-global zone is also associated with that non-global zone.

Any zone which is not the global zone is called a non-global zone. Most people call non-global zones simply "zones." Some people call them "local zones" but this is discouraged.

The default native zone file system model on Oracle Solaris 10 is called "sparse-root." This model emphasizes efficiency and security at the cost of some configuration flexibility. Sparse-root zones optimize physical memory and disk space usage by sharing some directories, like /usr and /lib. Sparse-root zones have their own private file areas for directories like /etc and /var. Whole-root zones increase configuration flexibility but increase resource usage. They do not use shared file systems for /usr, /lib, and a few others.

There is no supported way to convert an existing sparse-root zone to a whole-root zone. Creating a new zone is required.

NEW QUESTION 303

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:

rctladm - display and/or modify global state of system resource controls
The following command activates system logging of all violations of task.max-lwps.
rctladm -e syslog task.max-lwps
#

NEW QUESTION 304

The /etc/hosts file can be best described as .

- A. a local database of host names for rlogin, rsh, and rep
- B. the configuration file for the host name of the system
- C. a local database of information for the uname command
- D. the configuration file for the Domain Name Service (DNS)
- E. a local database of host names and their associated IP addresses

Answer: E

Explanation:

As your machine gets started, it will need to know the mapping of some hostnames to IP addresses before DNS can be referenced. This mapping is kept in the /etc/hosts file. In the absence of a name server, any network program on your system consults this file to determine the IP address that corresponds to a host name.

NEW QUESTION 307

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

```
ADDBOBJ      TYPE      STATE      ADDR
1o0/v4       static    ok         127-0.0.1/8
1o0/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 309

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
mirror-1	ONLINE	0	0	0
c3t5d0	ONLINE	0	0	0
c3t6d0	ONLINE	0	0	0

The following partition scheme is used for every disk drive in pool1:

```
ascii name = <ATA-VBOX HARDDISK-1.0-146.00GB>
bytes/sector = 512
sectors = 306184191
accessible sectors = 306184158
Part      Tag      Flag      First Sector      Size      Last Sector
0         usr      wm         256                145.99GB  306167774
1 unassigned wm         0                   0         0
2 unassigned wm         0                   0         0
3 unassigned wm         0                   0         0
4 unassigned wm         0                   0         0
5 unassigned wm         0                   0         0
6 unassigned wm         0                   0         0
8 reserved wm         306167775          8.00MB    306184158
```

Which two are true regarding the ZFS storage pool?

- A. The data on c3t3d0 is duplicated on c3t4do.
- B. The data is striped across disks c3t3d0 and c3t4do and mirrored across vdevs mirror-0 and mirror-1.
- C. The storage pool is 146 GB total size (rounded to the nearest GB).
- D. The storage pool is 584 G8 total size (rounded to the nearest GB).
- E. The storage pool is 292 GB total size (rounded to the nearest GB).

Answer: AE

NEW QUESTION 311

View the exhibit.

ADDROBJ	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

The configuration information in the exhibit is displayed on your system immediately after installing the OS. Choose the option that describes the selection made during the Installation of the OS to obtain this configuration.

- 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 DHCP network configuration option was chosen during the Installation of the OS.

Answer: A

Explanation:

There are two ways to configure the network configuration: automatic or manual. In the exhibit we see that DHCP has been used used. This indicates an automatic network configuration.

NEW QUESTION 316

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/osolBE 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 320

You are attempting to troubleshoot an event that should have made an entry into the messages log. This event happened about two weeks ago. Which file should you look at first?

- A. /var/adm/messages
- B. /var/adm/messages.0

- C. /var /adm/messagas.1
- D. /var/adm/messages.2
- E. /var/adm/messages.3

Answer: A

Explanation:

The /var/adm/messages is the file to which all the messages printed on the console are logged to by the Operating System. This helps to track back check the console messages to troubleshoot any issues on the system.

Syslog daemon also writes to this /var/adm/messages file.

The /var/adm/messages file monitored and managed by newsyslog and its configuration file is /usr/lib/newsyslog.

This script runs as the roots cron job everyday, checks the /var/adm/messages file and copies/moves it to /var/adm/messages.0, 1, 2, 3, 4, 5, 6, 7. In other words, it does the Log Rotation for the /var/adm/messages.

In an event the /var file system is running out of space, these files needs to checked and can be removed (not the actual /var/adm/messages itself) to free up space on the file system.

However, care has to be taken, if you decide to empty the /var/adm/messages itself for any reason. This process is called Truncation.

SOLARIS SYSTEM ADMIN TIPS, /var/adm/messages

NEW QUESTION 321

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 1Z0-821 Exam with Our Prep Materials Via below:

<https://www.certleader.com/1Z0-821-dumps.html>