

1Z0-821 Dumps

Oracle Solaris 11 System Administrator

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



NEW QUESTION 1

You are planning group names for a new system. You decide to use a numbering convention that includes the year and month the project began, to form the group number and name for work associated with that project.

So, for example, a project targeted to begin in January, 2013 would have the number (name):

201301(Pr201301)

What are the two problems with your plan?

- A. Group names may not contain a numeric character
- B. Group names may be no longer than 7 characters.
- C. Group numbers should not be larger than 60000.
- D. Group names should be all lowercase.

Answer: CD

Explanation:

C: The Group ID (GID) field contains the group's numerical ID. GIDs can be assigned whole numbers between 100 and 60000.

D: Group names contain only lowercase characters and numbers.

NEW QUESTION 2

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 3

You are troubleshooting network throughput on your server.

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

The correct command is .

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

Answer: B

Explanation:

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

NEW QUESTION 4

To confirm the IP addresses and netmasks have been correctly configured on the network interfaces, which command(s) should you use?

- A. ipadm show-if
- B. ipadm show-nic
- C. ipadm show-addr
- D. ipadm show-addr ipadm show-mask
- E. ipadm show-ip ipadm show-mask
- F. ipadm show-config

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.

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 5

The core dump configuration in your non-global zone is

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

A user is running a process in a non-global zone (testzone) and the process crashes. The process information is:
user126632618017:46:42pts/20:00/usr/bin/bash

When the user's process crashes in testzone, a non-global zone, where will the core dump be saved?

- A. The file will be stored in the non-global zone's directory: /var/core/pprocess/core.hash.2663.
- B. The file will be saved in the global zone's directory: /var/core/core.bash.2663.
- C. A core file cannot be generated in a non-global zone because it shares the kernel with the global zone.
- D. The file will be stored in the global zone's directory: /var/core/pprocess/core.bash.2663.
- E. The file will be saved in non-global zone's directory: /var/core/core.bash.2663

Answer: E

Explanation:

The line
init core file pattern: /var/core/core.%f.%p
will be used for the non-global process to determine the destination of the dump file.
Note: When a process is dumping core, up to three core files can be produced: one in the per-process location, one in the system-wide global location, and, if the process was running in a local (non-global) zone, one in the global location for the zone in which that process was running.

NEW QUESTION 6

To inspect network interface net3, you enter the following commands:

```
$ ipadm show-if | grep net3
net3      ip      down    no      --

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

What problem do you suspect? Assume the user is authorized and provided the correct password.

- A. The net3 interface hasn't been enabled yet.
- B. The net3 vnic hasn't been created.
- C. The net3/v4 ip object hasn't been configured.
- D. The net3 interface is not attached to a NIC or etherstub.

Answer: C

Explanation:

The following command marks the address object net1/v4a up that was previously marked down.
ipadm up-addr net1/v4a

NEW QUESTION 7

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 8

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

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

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

Answer: A

NEW QUESTION 9

Before booting testzone, a non-global zone, you want to connect to the zone's console so that you can watch the boot process. Choose the command used to connect to testzone's console.

- A. zoneadm – C testzone
- B. zoneadm – console testzone
- C. zlogin – z testzone console
- D. zlogin – z testzone – C
- E. zlogin – C testzone
- F. zoneadm – z testzone – C

Answer: E

NEW QUESTION 10

Select the five tasks that need to be performed on the Automated Installer (AI) install server before setting up the client.

- A. Create a local IPS repository on the AI Install server and start the repository server service, the publisher origin to the repository file.
- B. Set up a IP address on the AI install server.
- C. The DHCP server must be enabled on the install server and must provide the DHCP service for the clients.
- D. DHCP must be available on the network for the Install server and the clients, but the install server does not need to be the DHCP server.
- E. Download the AI boot image
- F. The image must be the same version as the Oracle Solaris OS that you plan to install on the client.
- G. Download the text install image into the IPS repository.
- H. Install the AI installation tools.
- I. Create the AI install service
- J. Specify the path to the AI network boot image ISO file and the path where the AI net image ISO file should be unpacked.
- K. Create the AI install service
- L. Specify the path to the AI network boot image ISO file and the path to the IPS repository.

Answer: BDFGI

Explanation:

B: Configure the AI install server to use a static IP address and default route.

D: The create-service command can set up DHCP on the AI install server. If you want to set up a separate DHCP server or configure an existing DHCP server for use with AI. The DHCP server must be able to provide DNS information to the systems to be installed.

E: An automated installation of a client over the network consists of the following high-level steps:

1. The client system boots over the network and gets its network configuration and the location of the install server from the DHCP server.
2. The install server provides a boot image to the client.
3. Characteristics of the client determine which installation instructions and which system configuration instructions are used to install the client.
4. The Oracle Solaris 11 OS is installed on the client, pulling packages from the package repository specified by the installation instructions in the AI install service.

G: Install the AI tool set.

Use the installadm create-service command to create an AI install service. Give the service a meaningful name, and specify the path where you want the service created. Specify the source of the network boot image (net image) package or ISO file.

```
installadm create-service [-n svcname] [-s FMRI_or_ISO] [-d imagepath]
```

```
-d imagepath
```

The imagepath is the location of the new install service. The install-image/solaris-auto- install package is installed to this location, or the specified ISO file is expanded at this location.

NEW QUESTION 10

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

```

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

```

```

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

```

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

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

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

Answer: C

NEW QUESTION 12

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 14

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 service.

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 16

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

A. bash shell

B. C shell

C. korn shell

D. bourne shell

Answer: A

Explanation:

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

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

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

* The legacy `ksh88` is available as `/usr/sunos/bin/ksh` from the `shell/ksh88` package.

* Korn shell compatibility information is available in `/usr/share/doc/ksh/COMPATIBILITY`.

NEW QUESTION 17

Examine this command and its output:

```
$ zfs list -r -t all tank
```

```
Name USED AVAIL REFER MOUNTPOINT
```

```
tank 3.00G 1.84G 32K /tank
```

```
tank/database 3.00G 1.84G 2.00G /tank/database tank/[email protected] 1.00G - 2.00G -
```

Which two conclusions can be drawn based on this output?

A. The tank dataset consumes 3 GB of storage.

B. The tank/ dataset consumes 1 GB of storage that is shared with its parent.

C. The tank/ dataset consumes 1 GB of storage that is not shared with its parent.

D. The tank/ dataset consumes 2 GB of storage that is shared with its child.

E. The tank/ dataset consumes 2 GB of storage that is not shared with its child.

Answer: AB

NEW QUESTION 21

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 25

A datalink can best be described as .

- A. a driver for a Network Interface Card
- B. the software connecting the Internet Layer and the Physical Layer
- C. a device that provides Classless Inter-Domain Routing
- D. a logical object used for IP Multipathing

Answer: D

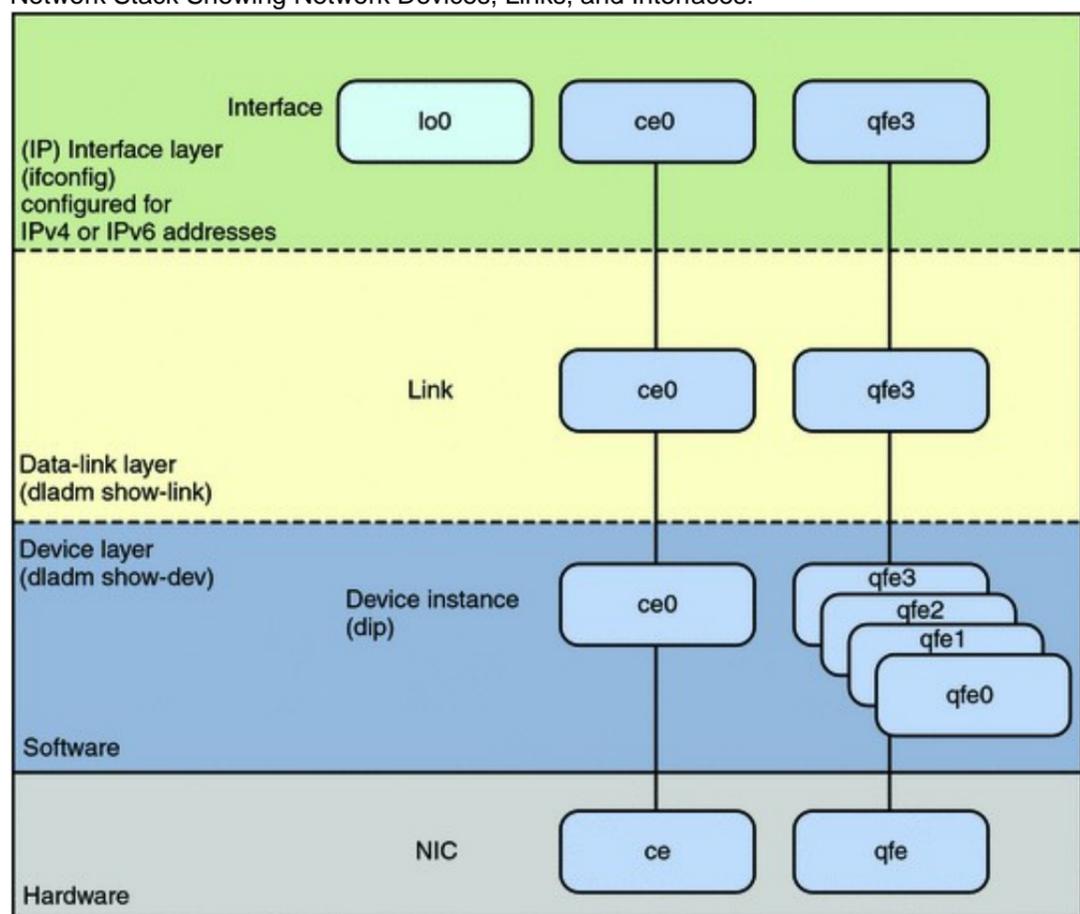
Explanation:

The command `dladm` is used to configure data-link interfaces in Sun Solaris. A configured data-link is represented in the system as interface that can be used for TCP/IP. Each data-link relies on either a single network device or an link aggregation device to send & receive packets.

Network interfaces provide the connection between the system and the network. These interfaces are configured over data links, which in turn correspond to instances of hardware devices in the system.

In the current model of the network stack, interfaces and links on the software layer build on the devices in the hardware layer. More specifically, a hardware device instance in the hardware layer has a corresponding link on the data-link layer and a configured interface on the interface layer. This one-to-one relationship among the network device, its data link, and the IP interface is illustrated in the figure that follows.

Network Stack Showing Network Devices, Links, and Interfaces:



NEW QUESTION 27

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 30

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 31

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 36

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 37

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 38

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 40

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 43

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 48

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 49

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:login <zonename> mpstat
- C. Run from the global zone:zonestar -r summary
- D. Run from the global zone:rctladm -1
- 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 51

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 52

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

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

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

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

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

Answer: A

NEW QUESTION 54

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

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

id_dsa id_dsa.pub id_rsa id_rsa.pub known_hosts authorized_keys Which two are true?

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

Answer: AE

Explanation:

A: You will see two files starting with id_rsa. id_rsa is the private key and id_rsa.pub is public key.

E: The .ssh/known_hosts file

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

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

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

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

```
newton 1024 35
153438062610297067329638677441205712613292203533062535600064224677647442
245028855505387934431717435134842994423656065076260604296084868001730665
553662299156116414854701274715680961503198280525759778667306417179500370
189017139564144825610347509023078143132936185076849630461827976942220442
313116255293297021841
ucsub 1024 37
132170811640421742212085598383135714069016332111955003414250071326834884
018721183646445780180633494496866895830879394309011412231102757022090299
732775466435482517698989962531081214859205054227533597152962802400251809
883548442498002326460312850336779152617243800769119880843882425555806081
435017335194477605333
simpson 1024 41
840896920592494584403453622735282634536002054701576247765078766974814128
393752943151071629834843909016027026612791643752972116459602750267266908
365259665072736159491719667576217171370458928680504368847255632477925660
234893185547218857655484574619075125368470792976275806263534208879722192
77539015703446529603
newton, 128.138.249.8 ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEA0d7Aoure0toNJ+YMYi61QP2ka8m5x5ZQIT7obP8C
K3eropfqsMPPY6uiylh9vpiFX2r1LHcbx139+vG6HOtVvuS8+lfMDtawm3WQvRuOopz3vVy
5GtMwtaOgehsXoT930Ryev1bH5myPtWklipITsOd2sX9k3tvjrmme4KCGGss=
```

NEW QUESTION 59

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 61

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 66

What is the result of executing the following command? `svcs -d svc:/network/ssh:default`

- A. disables the svc:/network/ssh:default service
- B. displays the services that svc:/network/ssh:default is dependent on
- C. displays the services that are dependent on the svc:/network/ssh:default service
- D. deletes the svc:/network/ssh:default service

Answer: B

Explanation:

The svcs command displays information about service instances as recorded in the service configuration repository.
-d Lists the services or service instances upon which the given service instances depend.

NEW QUESTION 71

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 72

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

```
~/bigscript &
```

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

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

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

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

Answer: A

Explanation:

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

grep: id 15166

ps -aef 15165

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

NEW QUESTION 76

Which two options are characteristics of a fast reboot?

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

Answer: CD

Explanation:

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

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

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

```
# svcadm refresh svc:/system/boot-config:default
```

D: Fast Reboot implements an in-kernel boot loader that loads the kernel into memory and then switches to that kernel.

The firmware and boot loader processes are bypassed, which enables the system to reboot within seconds.

The Fast Reboot feature is managed by SMF and implemented through a boot configuration service, svc:/system/boot-config. The boot-config service provides a

means for setting or changing the default boot configuration parameters. When the `config/fastreboot_default` property is set to true, the system performs a fast reboot automatically, without the need to use the `reboot -f` command. This property's value is set to true on the x86 platform. For task-related information, including how to change the default behavior of Fast Reboot on the SPARC platform, see [Accelerating the Reboot Process on an x86 Based System](#).

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

NEW QUESTION 79

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

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

Answer: E

Explanation:

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

NEW QUESTION 83

You have been asked to terminate a process that appears to be hung and will not terminate. The process table is shown below:
root 15163 15156 0 12:51:15 pts/3 0:00 hungscript What command will terminate the process?

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

Answer: A

Explanation:

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

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

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

Tip - When using the kill command to stop a process, first try using the command by itself, without including a signal option. Wait a few minutes to see if the process terminates before using the kill command with the -9 signal.

NEW QUESTION 84

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

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

Answer: C

Explanation:

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

The `fastreboot_default` property of the boot-config service enables an automatic fast reboot of the system when either the `reboot` or the `init 6` command is used.

When the `config/fastreboot_default` property is set to true the system automatically performs a fast reboot, without the need to use the `reboot -f` command. By default, this property's value is set to false on the SPARC platform and to true on the x86 platform.

NEW QUESTION 89

You run the command `dlstat show-link -r`.

Select the two correct statements regarding the information displayed in the INTRS column.

- A. No value is listed for virtual network interfaces.
- B. A value of 0 is listed for virtual interfaces and ether stubs.
- C. The number of Interrupts is listed, which indicates network efficiency.
- D. A number equal to the number of transmitted Ethernet frames is listed for physical links.
- E. The number of packets that were interrupted by a collision is listed, which may indicate hardware problems.

Answer: CE

Explanation:

In this output, the statistics for interrupt (INTRS) are significant. Low interrupt numbers indicate greater efficiency in performance. If the interrupt numbers are high, then you might need to add more resources to the specific link.

Example:

```
# dlstat -r -i 1
```

```
LINK IPKTS RBYTES INTRS POLLS CH<10 CH10-50 CH>50 e1000g0 101.91K 32.86M 87.56K 14.35K 3.70K 205 5  
nxge1 9.61M 14.47G 5.79M 3.82M 379.98K 85.66K 1.64K vnic1 8 336 0 0 0 0
```

e1000g0 0 0 0 0 0 0

nxge1 82.13K 123.69M 50.00K 32.13K 3.17K 724 24

vnic1 0 0 0 0 0 0

Note: `dlstat show-link [-r [-F] | -t] [-i interval] [-a] [-p] [-o field[, ...]] [-u R|K|M|G|T|P] [link]` Display statistics for a link.

-r
Display receive-side statistics only. Includes bytes and packets received, hardware and software drops, and so forth.

List of supported RX fields: link

usedby

ibytes ipkts intrs polls

hdrops: hardware drops

sdrops: software drops (owing to bandwidth enforcement) ch<10: number of packet chains of length < 10

ch10-50: number of packet chains of length between 10 and 50 ch>50: number of packet chains of length > 50

NEW QUESTION 91

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

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

Answer: A

Explanation:

Network configuration in Oracle Solaris 11 includes

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

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

NEW QUESTION 96

The line

`set noexec_user_stack=1`

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

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

Answer: D

Explanation:

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

Edit the `/etc/system` file, and add the following line: `set noexec_user_stack=1`

Reboot the system.

`# reboot`

NEW QUESTION 99

A local repository is available on this system and you need to enable clients to access this repository via HTTP. The repository information is:

PUBLISHERTYPESTATUSURI

solarisoriginonlinehttp://sysA.example.com

Identify two of the steps that are required to make the local repository on this server available to the client via HTTP.

- A. On the server: set the `pkg/inst_root` and `pkg/readonly` properties for the `svc:/application/pkg/server:default` service and enabled the service
- B. On the server: set the `sharefs` property on the ZFS file system containing the IPS repository.
- C. On the client: reset the origin for the solaris publisher.
- D. On the client: set the `pkg/inst_root` and `pkg/readonly` properties for the `svc:/application/server:default` service enable the service.
- E. On the client: start the `pkg.depotd` process.

Answer: AE

Explanation:

A: Configure the Repository Server Service

To enable clients to access the local repository via HTTP, enable the `application/pkg/server` Service Management Facility (SMF) service.

`# svccfg -s application/pkg/server setprop pkg/inst_root=/export/repoSolaris11`

`# svccfg -s application/pkg/server setprop pkg/readonly=true`

E: Use `pkg.depotd` to serve the repository to clients. Start the Repository Service

Restart the `pkg.depotd` repository service.

`# svcadm refresh application/pkg/server`

`# svcadm enable application/pkg/server`

To check whether the repository server is working, open a browser window on the localhost location.

NEW QUESTION 102

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

`whoami;who am i;id`

Which option correctly represents the output?

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

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

Answer: B

Explanation:

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

* who am i

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

```
user1 console Dec 30 20:20
```

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

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

Note:

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

Real UID

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

Effective UID

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

Saved UID

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

NEW QUESTION 104

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 109

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 114

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

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

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

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

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

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

- A) On Saturday:
zfs snapshot pool1/data@sat
zfs send pool1/data@sat > /remote/backup/full
On each weekday:
Remove the previous daily snapshot.
zfs snapshot pool1/data@daily
zfs send -i pool1/data@sat pool1/data@daily > /remote/backup/full
- B) On Saturday:
zfs create snapshot pool1/data@sat
zfs send pool1/data@sat |zfs recv remote/backup/`date +%m%d%y`
On each weekday:
Remove the previous daily snapshot.
zfs create pool1/data@daily
zfs send -i pool1/data@sat pool1/data@daily |zfs recv remote/backup/`date +%m%d%y`
- C) On Saturday:
zfs snapshot pool1/data@sat
zfs send pool1/data@sat > /remote/backup/full
On each weekday:
Remove the previous daily snapshot.
zfs snapshot pool1/data@daily
zfs send -i pool1/data@sat pool1/data@daily > /remote/backup/`date +%m%d%y`
- D) On Saturday:
zfs create snapshot pool1/data@sat
zfs send pool1/data@sat | zfs recv remote/backup
On each weekday:
Remove the previous daily snapshot.
zfs create -i pool1/data@sat pool1/data@daily
zfs send pool1/data@daily |zfs recv remote/backup

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

Answer: C

NEW QUESTION 115

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 deskto
- C. The GNOME desktop package must he 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 117

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 119

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

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

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

Answer: D

Explanation:

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

Boot Troubleshooting:

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

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

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

NEW QUESTION 122

Review the ZFS dataset output that is displayed on your system:

```
M  F  /data/file5
-  F  /data/file1
R  F  /data/file3 -> /data/file13
+  F  /data/file4
```

Which four correctly describe the output?

- A. /data/file4 has been added.
- B. The link /data/file3 has been added.
- C. /data/file3 has been renamed to /data/file13.
- D. /data/file4 has been modified and is now larger.
- E. /data/file1 has been deleted.
- F. /data/file1 has been modified and is now smaller.
- G. /data/file5 has been modified.
- H. /data/file3 (a link) has been removed.

Answer: ACEG

Explanation:

A: + Indicates the file/directory was added in the later dataset

C: R Indicates the file/directory was renamed in the later dataset E: - Indicates the file/directory was removed in the later dataset

G: M Indicates the file/directory was modified in the later dataset

Note: Identifying ZFS Snapshot Differences (zfs diff)

You can determine ZFS snapshot differences by using the zfs diff command.

The following table summarizes the file or directory changes that are identified by the zfs diff command.

File or Directory Change Identifier

* File or directory is modified or file or directory link changed M

* File or directory is present in the older snapshot but not in the newer snapshot

—

* File or directory is present in the newer snapshot but not in the older snapshot

+

* File or directory is renamed R

NEW QUESTION 125

After installing the OS, you boot the system and notice that the syslogd daemon is not accepting messages from remote systems.

Which two options should you select to modify the syslogd daemon configuration so that it accepts messages from remote systems?

A. svccfg -s svc:/system/system -log setprop start/exec= "syslogd -t"Restart the syslogd daemon.

B. Set the following parameter in the /etc/syslogd.conf file: LOG_FROM_REMOTE= YESRestart the syslogd daemon.

C. svcadm enable svc:/system/system -log/config/log_from_remoteRestart the syslogd daemon.

D. svccfg -s svc:/system/system-log setprop config/log_from_remote=trueRestart the syslogd daemon.

E. Set the following parameter in the /etc/default/syslogd file: LOG_FROM_REMOTE=YESRestart the syslogd daemon.

Answer: BD

Explanation:

B: The /etc/default/syslogd file contains the following default parameter settings. See FILES.

LOG_FROM_REMOTE

Specifies whether remote messages are logged. LOG_FROM_REMOTE=NO is equivalent to the -t command-line option. The default value for LOG_FROM_REMOTE is YES.

NEW QUESTION 129

Examine this command and its output:

```
# zoneadm list Global
```

```
zone2
```

Which two outcomes can be deduced from this output?

A. There is exactly one nonglobal zone installed.

B. There is one nonglobal zone running.

C. There is at least one oneglobal zone configured.

D. There is one nonglobal zone that is not configured.

E. There is one nonglobal zone that is not running.

F. There is one nonglobal zone that is not installed.

G. There is exactly one nonglobal zone configured.

Answer: BC

NEW QUESTION 134

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 138

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

A. Jumpstart

B. LiveCD

C. A text Installation over the network

D. An Automated Installation performed on an x86 client

E. An Automated Installation using media from a local DVD or USB drive

F. An Automated Installation using a networked repository

Answer: DEF

Explanation:

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

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

NEW QUESTION 143

When you issue the “gzip: zommand not found” message is displayed. You need to install the gzip utility on your system. Which command would you use to check if the gzip utility is available from the default publisher for installation?

- A. pkg info|grep gzip
- B. pkg list SUNWgzip
- C. pkg contents gzip
- D. pkg search gzip

Answer: D

Explanation:

Searching for Packages

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

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

pkg search

search [-H|a|f|p|r] [-o attribute ...] [-s repo_uri] query Search for matches to the query, and display the results.

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

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

NEW QUESTION 144

You want the system to generate an email notification each time one of the services has changed its state. Which option would send an email message to the system administrator whenever a service changes to the maintenance state?

- A. Use the setsc command in ALOM to enable the mail alerts to be sent to a specified email address whenever the fault management facility detects a service change to the maintenance state.
- B. Make an entry in the /etc/syslog.conf file to instruct syslogd to send an email alert when it receives a message from the SMF facility that a service has changed to the maintenance state.
- C. Use the svccfg setnotify command to create a notification and send an email when a service enters the maintenance state.
- D. Use the scvadm command to enable the notification servic
- E. Set the -g maintenance option on the netnotify service to send an email when a service enters the maintenance state.

Answer: C

Explanation:

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:

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
```

NEW QUESTION 145

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 146

Which three statements accurately describe the Automated Installation (AI) client?

- A. If the AI client does not match any criteria to use a custom manifest or script, the default manifest is used.

- B. If the AI client does not match any criteria to use a custom manifest or script, the automated installation aborts.
- C. Any manifest or script in a service can be designated to be the default for that service.
- D. Only the default.xml file is used as the default AT client manifest.
- E. If a client system does not use any SC profile, then an interactive tool opens on that client at first boot after that client installation to complete the configuration of that client.
- F. If a client system does not use any SC profile, then the install server will use the default SC profile.

Answer: ACE

Explanation:

Each client uses one and only one AI manifest to complete its installation. The AI manifest is selected for a client according to the following algorithm:

* If no custom AI manifests are defined for this install service, the default AI manifest is used. The default AI manifest is not associated with any client criteria etc. Each client can use any number of system configuration profiles. If a client system does not use any configuration profile, then an interactive tool opens on that client at first boot after that client installation to complete the configuration of that client.

NEW QUESTION 150

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 152

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 153

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

```
jack@solaris:/export/home/jack $ ls File filea Filea fileb Fileb filec Filec
```

Which two statements are correct?

- A. The pattern [?i]*a will expand to filea Filea.
- B. The pattern [fF]*a? will expand to [fF] *a?.
- C. The pattern [gfe] * will expand to file filea fileb filec.
- D. The pattern [g-e] * will expand to file filea fileb filec.
- E. The pattern [fF] [a-zA-z] i*e will expand to file.

Answer: AC

Explanation:

A: starting with one single character, second character must be letter i, any characters, ending with letter a.

C: starting with letter e, f, or g, followed by anything.

NEW QUESTION 157

You are having an issue with the shutdown command. You wish to determine if the file is a script or an executable program. Which command would you use to determine this?

- A. od shutdown

- B. file shutdown
- C. test shutdown
- D. cksum shutdown
- E. attrib shutdown

Answer: B

Explanation:

The file command determines the file type file tests each argument in an attempt to classify it. There are three sets of tests, performed in this order: filesystem tests, magic tests, and language tests. The first test that succeeds causes the file type to be printed.

NEW QUESTION 158

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

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

Remote system name: backupSYS

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

Most recent backup name: rpool-1202

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

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

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

Answer: A

Explanation:

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

* ZFS root pool cannot be recovered

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

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

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

Boot from CD/DVD or the network.

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

ok boot net -s

ok boot cdrom -s

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

2.

Mount the remote snapshot dataset. For example:

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

3.

Recreate the root pool. For example:

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

4.

Restore the root pool snapshots.

This step might take some time. For example:

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

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

5.

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

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

Install the boot blocks on the new disk.

On a SPARC based system:

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

NEW QUESTION 163

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

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

Answer: B

Explanation:

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

```
# ipadm show-addr
```

```
ADDROBJ TYPE STATE ADDR
```

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

NEW QUESTION 167

Select two correct statements about the authentication services available in Oracle Solaris 11.

- A. Pluggable Authentication Modules (PAM) is used to control the operation of services such console logins and ftp.
- B. The Secure Shell can be configured to allow logins across a network to remote servers without transmitting passwords across the network.
- C. Secure Remote Procedure Calls (Secure RPC) provides a mechanism to encrypt data on any IP Socket connection.
- D. Pluggable Authentication Modules (PAM) is used to implement the Secure Shell in Oracle Solaris 11.
- E. Simple Authentication and Security Layer (SASL) provides a mechanism to authenticate and encrypt access to local file system data.

Answer: AE

Explanation:

A: Pluggable Authentication Modules (PAM) are an integral part of the authentication mechanism for the Solaris. PAM provides system administrators with the ability and flexibility to choose any authentication service available on a system to perform end-user authentication. By using PAM, applications can perform authentication regardless of what authentication method is defined by the system administrator for the given client. PAM enables system administrators to deploy the appropriate authentication mechanism for each service throughout the network. System administrators can also select one or multiple authentication technologies without modifying applications or utilities. PAM insulates application developers from evolutionary improvements to authentication technologies, while at the same time allowing deployed applications to use those improvements. PAM employs run-time pluggable modules to provide authentication for system entry services.

E: The Simple Authentication and Security Layer (SASL) is a method for adding authentication support to connection-based protocols. Simple Authentication and Security Layer (SASL) is a framework for authentication and data security in Internet protocols. It decouples authentication mechanisms from application protocols, in theory allowing any authentication mechanism supported by SASL to be used in any application protocol that uses SASL. Authentication mechanisms can also support proxy authorization, a facility allowing one user to assume the identity of another. They can also provide a data security layer offering data integrity and data confidentiality services. DIGEST-MD5 provides an example of mechanisms which can provide a data-security layer. Application protocols that support SASL typically also support Transport Layer Security (TLS) to complement the services offered by SASL.

NEW QUESTION 168

Which two options accurately describe the network characteristics of a zone?

- A. DHCP address assignment cannot be configured in a shared IP zone.
- B. Shared IP is the default type of network configuration.
- C. Exclusive IP is the default type of network configuration.
- D. By default, all IP addresses, netmasks, and routes are set by the global zone and cannot be altered in a non global zone.
- E. IPMP cannot be managed within the non-global zone.
- F. Commands such as snoop and dladm cannot be used on datalinks that are in use by a running zone.

Answer: AB

Explanation:

A: Non-global zones can not utilize DHCP (neither client nor server).

B (not C): By default, non-global zones will be configured with a shared IP functionality. What this means is that IP layer configuration and state is shared between the zone you're creating and the global zone. This usually implies both zones being on the same IP subnet for each given NIC.

Note: A zone is a virtual operating system abstraction that provides a protected environment in which applications run. The applications are protected from each other to provide software fault isolation. To ease the labor of managing multiple applications and their environments, they co-exist within one operating system instance, and are usually managed as one entity.

The original operating environment, before any zones are created, is also called the "global zone" to distinguish it from non-global zones, The global zone is the operating system instance.

Incorrect Answer

E: Exclusive-IP zones can use IPMP. IPMP is configured the same way in an exclusive-IP zone as it is on a system not using zones. For shared-IP zones, IPMP can be configured in the global zone. F: Full IP-level functionality is available in an exclusive-IP zone. An exclusive-IP zone has its own IP-related state.

An exclusive-IP zone is assigned its own set of data-links using the zonecfg command. The zone is given a data-link name such as xge0, e1000g1, or bge32001, using the physical property of the net resource. The address property of the net resource is not set.

Note that the assigned data-link enables the snoop command to be used.

The dladm command can be used with the show-linkprop subcommand to show the assignment of data-links to running exclusive-IP zones.

NEW QUESTION 169

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

crontab: you are not authorized to use cron. Sorry

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 172

.....

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>