

## 1Z0-821 Dumps

### Oracle Solaris 11 System Administrator

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



### NEW QUESTION 1

You are the administrator for a group of shell script developers. They use vi, and have asked you to make their scripts automatically executable when they save their files.

How can this be accomplished?

- A. Enter set -o vi on the command line, or include it in each user's startup script.
- B. Enter umask -s on the command line, or include it in each user's startup script.
- C. Enter umask 000 on the command line, or include it in each user's startup script.
- D. Enter umask 777 on the command line, or include it in each user's startup script.
- E. It is not possible to automatically set the execute bit on with the umask setting, or vi option.
- F. Enter umask 766 the command line, or include it in the global startup script for the default shell.

**Answer:** E

**Explanation:** Unlike DOS, which uses the file extension to determine if a file is executable or not, UNIX relies on file permissions.

The value assigned by umask is subtracted from the default.

User's file creation mask. umask sets an environment variable which automatically sets file permissions on newly created files. i.e. it will set the shell process's file creation mask to mode.

umask 000 would grant full permissions. Note: 777 full permissions

### NEW QUESTION 2

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 3

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

NAME	USED	AVAIL	REFER	MOUNTPPOINT
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 4

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 5

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 6

Review the zonestat command:

zonestate - q physical - memory -R high -z -p -p "zones" 10 24h 60m

Select the option that correctly describes the information that is displayed by this command.

- A. It is a sample of dbzone's physical memory usage taken every hour over a 24-hour period. Only the top 10 samplings of peak memory usage are displayed.
- B. All other utilization data is eliminated.
- C. It is a sample of dbzone's CPU, virtual memory, and networking utilization. Physical memory is executed from the report. The sampling is taken every 10 minutes over a 24-hour period and peak utilization is displayed each hour.
- D. It is a sample of dbzone's CPU, virtual memory, and networking utilization. Physical memory is executed from the report. The sampling is taken every 10 minutes over a 24-hour period and displayed each hour.
- E. It is a sample of dbzone's physical memory usage taken every 10 seconds and 24-hour period. Only peak virtual memory usage and CPU utilization are displayed each hour. All other Utilization data is eliminated.
- F. It is a sample of dbzone's physical memory usage taken every 10 seconds and 24-hour period. Only peak memory usage is displayed each hour. All other utilization data is eliminated.

**Answer: D**

**Explanation:** \* (Not A, B, C): interval (here 10 seconds): Specifies the length in seconds to pause between each interval report.

\* duration (here 24 h)

\* -R report[, report] (here high) Print a summary report.

High Print a summary report detailing the highest usage of each resource and zone during any interval of the zonestat utility invocation.

Note: The zonestat utility reports on the cpu, memory, and resource control utilization of the currently running zones. Each zone's utilization is reported both as a percentage of system resources and the zone's configured limits.

The zonestat utility prints a series of interval reports at the specified interval. It optionally also prints one or more summary reports at a specified interval.

The default output is a summary of cpu, physical, and virtual memory utilization. The -r option can be used to choose detailed output for specific resources.

#### NEW QUESTION 7

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

Select two unsupported items that will require changes.

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

**Answer: AD**

**Explanation:** A: terminal

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

D: system\_locale

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

#### NEW QUESTION 8

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 9

When speaking to an Oracle Support Engineer, you are asked to verify the version of the Solaris 11 build currently running on your system.

Which command would display the Solaris 11 build version currently running on your system?

- A. pkg info all
- B. cat /etc/release
- C. cat /etc/update
- D. prtconf | grep -i update
- E. pkg info entire

**Answer: B**

**Explanation:** Which Solaris release you are running on your system can be determined using the following command:

cat /etc/release

This will tell you which release you are running and when it was released. The more recent your system, the more info is contained in this file.

Example:

# cat /etc/release

Oracle Solaris 10 8/11 s10s\_u10wos\_17b SPARC

Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved. Assembled 23 August 2011

#### NEW QUESTION 10

Which modification needs to be made to the Service Management Facility before you publish a new package to the IPS repository?



- A. The pkg.depotd must be disabled.
- B. The pkg/readonly property for the application/pkg/server service must be set to false.
- C. The Pkg/writable\_root property for the application/Pkg/server service must be set to true.
- D. The pkg/image.root property for the application/pkg/server service must be set to the location of the repository.

**Answer:** D

**Explanation:** pkg/image\_root

(astring) The path to the image whose file information will be used as a cache for file data.

#### NEW QUESTION 10

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 13

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

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

**Answer:** BDE

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

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

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

#### NEW QUESTION 15

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

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

You then ran: `ipadm show-if`

The result indicated that the interface was down.

You then ran:

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

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

The last command indicated that the interface was up.

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

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

**Answer:** B

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

#### NEW QUESTION 17

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 21

On server A, you enter the following command to add a static route to serverA route -p add

```
-host 192.168.1.101 192.168.1.101 -static
```

What is the purpose of this command?

- A. to temporarily bypass IP Filter rules
- B. to specify an IPMP target IP address to in.mpathd
- C. to specify routing to an adjacent network when in.rdisc is not used
- D. to specify routing to an adjacent network when in.routed is not used
- E. to ensure the IP address for serverB is not flushed from the ARP cache
- F. to optimize link aggregation using a direct connection between two systems

**Answer:** B

**Explanation:** Note: # route -p add -host destination-IP gateway-IP -static

where destination-IP and gateway-IP are IPv4 addresses of the host to be used as a target.

For example, you would type the following to specify the target system 192.168.10.137, which is on the same subnet as the interfaces in IPMP group itops0:

```
$ route -p add -host 192.168.10.137 192.168.10.137 -static
```

This new route will be automatically configured every time the system is restarted. If you

want to define only a temporary route to a target system for probe-based failure detection, then do not use the -p option.

#### NEW QUESTION 24

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 28

When setting up Automated Installer (AI) clients, an interactive tool can be used to generate a custom system configuration profile. The profile will specify the time zone, data and time, user and root accounts, and name services used for an AI client installation. This interactive tool will prompt you to enter the client information and an SC profile (XML) will be created.

Which interactive tool can be used to generate this question configuration?

- A. sys-unconfig
- B. installadm set-criteria
- C. sysconfig create-profile
- D. installadm create-profile

**Answer:** B

**Explanation:** Use the installadm set-criteria command to update the client criteria associated with an AI manifest that you already added to a service using installadm add-manifest.

Use the installadm add-manifest command to add a custom AI manifest to an install service.

The value of manifest is a full path and file name with .xml extension. The manifest file contains an AI manifest (installation instructions). The manifest file can also

reference or embed an SC manifest (system configuration instructions).

**NEW QUESTION 29**

Review the boot environment information displayed on your system:

oldBE	-	-	149.OK	static	2011-11-28	15:15
newBE	!	-	363.05M	static	2011-11-28	14:47
solaris	-	-	100.68M	static	2011-11-20	18:09
solaris-1	NR	/	19.07G	static	2012-01-22	07:23

Which two options accurately describe the newBE boot environment?

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

**Answer:** BC

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

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

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

You cannot create a snapshot of an unbootable boot environment.

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

You cannot rename an unbootable boot environment. (C)

**NEW QUESTION 33**

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 37**

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 41**

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

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

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

**Answer:** D

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

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

How to Create a Boot Environment

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

# beadm create BeName

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

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

# beadm mount BeName mount-point

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

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

4. (Optional) Activate the boot environment.

# beadm activate BeName

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

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

#### NEW QUESTION 43

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 46

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

iusedby

ibytes ipkts intrs polls

hdrops: hardware drops

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

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

#### NEW QUESTION 51

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



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

Which two options describe the solaris-1 boot environment?

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

**Answer:** AE

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

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

N – Active now.

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

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

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

## NEW QUESTION 52

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 57

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 59

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 62

Consider the following commands:

```
rm file1
echo "Hello, world" > file2
cat file1 && cat file2
```

What is displayed when this sequence of commands is executed using the bash shell?

- A. Hello, world
- B. cat: cannot open file1
- C. cat: cannot open file1Hello, world
- D. cat: cannot open file1 Hello, World
- E. bash: syntax error near unexpected token '&&'

**Answer:** B

**Explanation:** First line (rm file1) deletes/removes file1. Second line captures the text into file2.

The first part of line 3 (cat file1) fails as the file1 does not exist.

The && (AND) operator will ensure that the third line fails. The result of line 3 will be the result of first part of line 3 (cat file1).

Note: cat - concatenate files and print on the standard output

Note #1: A list is a sequence of one or more pipelines separated by one of the operators ';', '&', '&&', or '||', and optionally terminated by one of ';', '&', or a newline.

Of these list operators, '&&' and '||' have equal precedence, followed by ';' and '&', which have equal precedence.

AND and OR lists are sequences of one or more pipelines separated by the control operators '&&' and '||', respectively. AND and OR lists are executed with left associativity.

An AND list has the form command1 && command2

command2 is executed if, and only if, command1 returns an exit status of zero.

An OR list has the form command1 || command2

command2 is executed if, and only if, command1 returns a non-zero exit status.

The return status of AND and OR lists is the exit status of the last command executed in the list.

Note #2 (on exit status): Zero means command executed successfully, if exit status returns non-zero value then your command failed to execute.

#### NEW QUESTION 64

The ZFS configuration on your server is:

```
Pool1 6.67G31K/pool Pool1/data31K31K/data
```

Select the three commands that you would use to 1. Create, 2. List, and 3. Delete a snapshot of the /data file system.

- A. zfs snapshot pool1/data@now
- B. zfs create snapshot pool1/data@now
- C. zfs list -t snapshot
- D. zfs list -t snapshot pool1/data
- E. zfs destroy pool1/data@now
- F. zfs destroy snapshot pool1/data@now

**Answer:** ADE

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

D: You can list snapshots as follows:

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

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

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

#### NEW QUESTION 65

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 70

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 75

A change in your company's security policy now requires an audit trail of all administrators assuming the sysadm role, capturing:

There are two command necessary to accomplish this change. One is a rolemod command. What is the other?

- A. auditconfig set policy=argv
- B. auditconfig -setpolicy +argv
- C. auditconfig -setflags lo, ex sysadm
- D. auditconfig set flags=lo, ex sysadm

**Answer:** B

**Explanation:** Audit Significant Events in Addition to Login/Logout (see step 2 below)

Use this procedure to audit administrative commands, attempts to invade the system, and other significant events as specified by your site security policy.

For all users and roles, add the AUE\_PFEEXEC audit event to their preselection mask.

```
# usermod -K audit_flags=lo, ps:no username
```

```
# rolemod -K audit_flags=lo, ps:no rolename
```

```
# auditconfig -setpolicy +argv
```

3- Record the environment in which audited commands are executed.

```
# auditconfig -setpolicy +arge
```

Note: [-t] -setpolicy [+|-]policy\_flag[, policy\_flag ...]

Set the kernel audit policy. A policy policy\_flag is literal strings that denotes an audit policy. A prefix of + adds the policies specified to the current audit policies. A prefix of - removes the policies specified from the current audit policies. No policies can be set from a local zone unless the perzone policy is first set from the global zone.

#### NEW QUESTION 80

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 85

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 87

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/[email protected] snapshot depends on the tank/database dataset.
- D. It fails because the tank/[email protected] clone depends on the tank/database dataset.
- E. It fails because the tank/database data set is not empty.

**Answer:** C

#### NEW QUESTION 89

You create a flash archive of the Solaris 10 global zone on the serves 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/bactup\_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. zone cfg - 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 94

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 98

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

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

**Answer:** D

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

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

#### NEW QUESTION 101

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+IfMDtawm3WQvRuOopz3vVy
5GtMwtaOgehsXoT930Ryev1bH5myPtWKlipITsOd2sX9k3tvjrmme4KCGGss=
```

#### NEW QUESTION 102

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

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

F. Default minimum length is 6, and is defined in /usr/sadm/defadduser.

**Answer:** B

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

#### NEW QUESTION 104

Which network protocol is responsible for routing packets from one network to another?

- A. TCP
- B. UDP
- C. IP
- D. ICMP
- E. Ethernet

**Answer:** C

**Explanation:** The Internet Protocol (IP) is the principal communications protocol in the Internet protocol suite for relaying datagrams across network boundaries. Its routing function enables internetworking, and essentially establishes the Internet.

#### NEW QUESTION 109

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 114

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

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

**Answer:** B

**Explanation:** The User Datagram Protocol (UDP) is one of the core members of the Internet Protocol Suite, the set of network protocols used for the Internet. With

UDP, computer applications can send messages, in this case referred to as datagrams, to other hosts on an Internet Protocol (IP) network without requiring prior communications to set up special transmission channels or data paths. Compared to TCP, UDP is a simpler message-based connectionless protocol.

**NEW QUESTION 117**

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 122**

In a fresh installation of Oracle Solaris 11, default datalinks are named with a genetic naming convention, and they increment as you add interfaces. What is the default name?

- A. eth#
- B. net#
- C. el000g#
- D. lo#
- E. nic#

**Answer:** B

**Explanation:** When you install this Oracle Solaris release on a system for the first time, Oracle Solaris automatically provides generic link names for all the system's physical network devices. This name assignment uses the net# naming convention, where the # is the instance number. This instance number increments for each device, for example, net0, net1, net2, and so on.

Note:

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

**NEW QUESTION 123**

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 127**

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 132**

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 137

Which command should you choose to display the current parameters for the FSS scheduler?

- A. dispadmin -c FSS
- B. prionctl -c FSS
- C. dispadmin -c FSS -g
- D. priocntl -c FSS -g

**Answer:** C

**Explanation:** The dispadmin command displays or changes process scheduler parameters while the system is running.

-c class  
Specifies the class whose parameters are to be displayed or changed. Valid class values are: RT for the real-time class, TS for the time-sharing class, IA for the inter-active class, FSS for the fair-share class, and FX for the fixed-priority class. The time-sharing and inter-active classes share the same scheduler, so changes to the scheduling parameters of one will change those of the other.

-g  
Gets the parameters for the specified class and writes them to the standard output.

#### NEW QUESTION 139

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 144

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 149

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 154

Identify three differences between the shutdown and init commands.

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

**Answer:** ABE

#### NEW QUESTION 159

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 163

To help with your troubleshooting, you need to determine the version of the OBP. Which two commands will provide you with this information?

- A. printenv
- B. banner
- C. .version
- D. set-env
- E. show-devs
- F. value version

**Answer:** BC

**Explanation:** B: banner

Displays power-on banner.

The PROM displays the system banner. The following example shows a SPARCstation 2 banner. The banner for your SPARC system may be different.

SPARCstation 2, Type 4 Keyboard

ROM Rev. 2.0, 16MB memory installed, Serial # 289 Ethernet address 8:0:20:d:e2:7b, Host ID: 55000121 C: .version

Displays version and date of the boot PROM.

Note: OBP-OpenBootProm is a firmware which is placed on the sun machine's prom chip. It is a os independent user interface to deal with the sun machine's hardware components. The user interface provides one or more commands to display system information.

**NEW QUESTION 165**

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 170**

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 172**

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 175**

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 177**

Which two options describe how to override the default boot behavior of an Oracle Solaris 11 SPARC system to boot the system to the single-user milestone?

- A. from the ok prompt, issue this command: boot -m milestone=single-user
- B. From the ok prompt, issue this command: boot -m milestone/single-user
- C. From the ok prompt, issue this command: boot -milestone=single-user
- D. From the ok prompt
- E. issue this command:boot -s
- F. From from the ok prompt, issue this command:boot -m milestone=s

**Answer:** AD

**Explanation:** By default, Solaris will boot to the pseudo milestone “all” and start all services. This behaviour can be changed at boot time using either “-s” to reach single-user, or the new SMF option “-m milestone=XXX” (see kernel(1M) for a list of the bootable milestones) to select an explicit milestone.

Note: boot -s is the same as: boot -m milestone=single-user

with the difference being that the former is a lot less to type and is what most SysAdmins will be familiar with.

#### NEW QUESTION 179

Your mentor suggests using the dladm rename-link command to rename the network datalinks.

What are the two advantages of following this advice?

- A. It can clarify which network interface has what purpose.
- B. It can simplify specifying the network interface with the dladm modify-aggr command.
- C. It can simplify specifying the network interface with the dladm modify-bridge command.
- D. It can simplify IP filter rule changes if the network interface is replaced with a different type.
- E. It can prevent accidental deletion of the network interface with the dladm delete-phys command.
- F. It can prevent accidental deletion of the network interface configuration with the ipadm delete-addr command.

**Answer:** AD

**Explanation:** Note: dladm rename-link [-R root-dir] link new-link

Rename link to new-link. This is used to give a link a meaningful name, or to associate existing link configuration such as link properties of a removed device with a new device.

#### NEW QUESTION 182

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. Identify two requirements for using AI.

- A. You should set up DHC
- B. The DHCP server and AI install server can be the same machine or two different machines.
- C. You can create only one manifest per install servic
- D. If you need more than one manifest, you should create multiple install services.
- E. The minimum requirement to operate as an AI install server is 1 GB of memory.
- F. If two client machines have different architectures and need to be installed with the same version of the Oracle Solaris 11 OS, then you should create two AI manifests and a single install service.

**Answer:** AD

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

Step 1. The client system boots over the network and gets its network configuration and the location of the install server from the DHCP server.

Step 2: The install server provides a boot image to the client. Etc.

D: 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.

#### NEW QUESTION 187

\_\_\_\_\_ serves as the interface between the SMF repository and the user to ensure that a consistent, picture of the repository is presented to the user.

- A. repository.db
- B. service manifest
- C. svc.startd
- D. svc.configd

**Answer:** D

**Explanation:** SVC.CONFIGD is the repository daemon responsible for maintaining /etc/svc/repository.db. The repository.db must come clean during this integrity check otherwise it is a "no go" for usual boot sequence to run level 3. The repository may get corrupted due to various hardware issues, software bugs, disk write failures, etc.

Note: When svc.configd(1M), the Solaris Repository Daemon, is started, it does an integrity check of the smf(5) repository, stored in /etc/svc/repository.db. This integrity check can fail due to a disk failure, the database file being corrupted either due to a hardware bug, a software bug, or an accidental overwrite. If the integrity check fails, svc.configd will write a message to the console.

#### NEW QUESTION 188

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

- A. Updates are only available to customers with an active support contrac
- B. The updates are distributed through the My Oracle Support web portal and installed in a central locatio
- C. All software packages are then updated manually from the command line using the smpatch command.
- D. Patches are download from <http://support.oracle.com> either automatically or manuall
- 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 repositor
- 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 imag
- 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 192**

You are troubleshooting the failure of a computer to mount an NFS file system hosted by a server (hostname mars) in the local area network. Select the three commands that will enable you to identify the problem.

- A. ping -s mars
- B. cat /etc/vfstab
- C. cat /etc/dfs/dfstab
- D. sharemgr show -v
- E. showmount -e mars
- F. rpcinfo -s mars | egrep 'nfs|mountd'

**Answer:** BEF

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

Mount: mount-point /DS9 does not exist.

To solve the mount point error condition, check that the mount point exists on the client. Check the spelling of the mount point on the command line or in the /etc/vfstab file (B) on the client, or comment out the entry and reboot the system.

Note: The /etc/vfstab file lists all the file systems to be automatically mounted at system boot time, with the exception of the /etc/mnttab and /var/run file systems.

E: showmount

This command displays all clients that have remotely mounted file systems that are shared from an NFS server, or only the file systems that are mounted by clients, or the shared file systems with the client access information. The command syntax is:

showmount [ -ade ] [ hostname ]

where -a prints a list of all the remote mounts (each entry includes the client name and the

directory), -d prints a list of the directories that are remotely mounted by clients, -e prints a list of the files shared (or exported), and hostname selects the NFS server to gather the information from. If hostname is not specified the local host is queried.

F: \* mountd Daemon

This daemon handles file-system mount requests from remote systems and provides access control. The mountd daemon checks /etc/dfs/sharetab to determine which file systems are available for remote mounting and which systems are allowed to do the remote mounting.

\* Commands for Troubleshooting NFS Problems

These commands can be useful when troubleshooting NFS problems. rpcinfo Command

This command generates information about the RPC service that is running on a system.

**NEW QUESTION 193**

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 196**

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 197**

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 201

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 stripped across four disk components.
- B. The rsdz1-0 and c3t640 components are submirrors of pool1.
- C. Data will only be stripped 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 203

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

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

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

**Answer:** A

#### NEW QUESTION 207

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 210

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 212

User jack on host solaris attempts to use ssh to log in to host oracle and receives this message:

```
jack@solaris:~$ ssh oracle
```

```
ssh: connect to host oracle port 22: connection refused What is the problem?
```

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

**Answer:** G

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

#### NEW QUESTION 216

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 217

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 219

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 220

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 222

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 225

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 227

A user account must be a member of a primary group, and may also be a member of one or more secondary groups. What is the maximum total number of groups that one user can concurrently belong to?

- A. 15
- B. 16
- C. 17
- D. 63
- E. 64
- F. 65
- G. The number of groups one user can concurrently belong to is unlimited in Solaris 11.

**Answer:** B

**Explanation:** Each user belongs to a group that is referred to as the user's primary group. The GID number, located in the user's account entry within the /etc/passwd file, specifies the user's primary group.

Each user can also belong to up to 15 additional groups, known as secondary groups. In the /etc/group file, you can add users to group entries, thus establishing the user's secondary group affiliations.

Note (4 PSARC/2009/542):

his project proposes changing the maximum value for NGROUPS\_MAX from 32 to 1024 by changing the definition of NGROUPS\_UMAX from 32 to 1024.

The use for a larger number of groups is described in CR 4088757, particular in the case of Samba servers and ADS clients; the Samba servers map every SID to a Unix group. Users with more than 32 groups SIDs are common. We've seen reports varying from "64 is enough", "128 is absolutely enough" and "we've users with more 190 group SIDS).

NGROUPS\_MAX as defined by different Unix versions are as follows (<http://www.j3e.de/ngroups.html>):

Linux Kernel >= 2.6.3 65536

Linux Kernel < 2.6.3 32 Tru64 / OSF/1 32

IBM AIX 5.2 64

IBM AIX 5.3 ... 6.1 128

OpenBSD, NetBSD, FreeBSD, Darwin (Mac OS X) 16 Sun Solaris 7, 8, 9, 10 16 (can vary from 0-32)

HP-UX 20

IRIX 16 (can vary from 0-32)

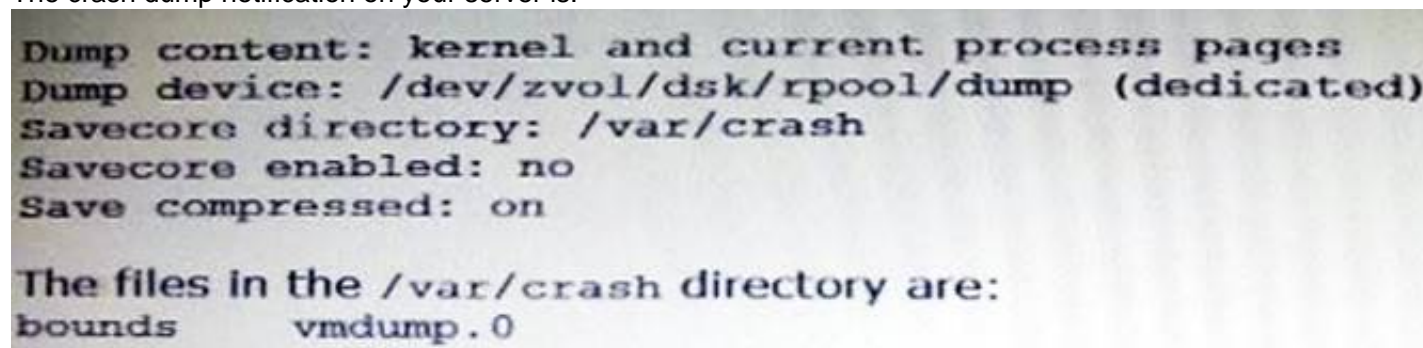
Plan 9 from Bell Labs 32

Minix 3 0 (Minix-vmd: 16)

QNX 6.4 8

#### NEW QUESTION 232

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 theses 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 235

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 239

To assist in examining and debugging running processes, Solaris 11 has a utility that returns process 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 243

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 244

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 249

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

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

- A. pkgchk opendap
- B. pkginfo opendap
- C. pkg contents opendap
- D. pkg verify opendap
- 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 252

View the Exhibit and review the disk configuration.

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

Primary label contents:

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

format> █
```

The following command is executed on the disk: `zpool create pool1 c3t3d0s0`  
What is the result of executing this command?

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

**Answer: B**

## NEW QUESTION 254

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

```
pool: pool1
state: ONLINE
scan: none requested
config:

    NAME                STATE                READ WRITE CKSUM
    pool1                ONLINE                0     0     0
      mirror-0           ONLINE                0     0     0
        c4t0d0           ONLINE                -     -     -
        c4t1d0           ONLINE                -     -     -
      mirror-1           ONLINE                0     0     0
        c4t2d0           ONLINE                -     -     -
        c4t3d0           ONLINE                -     -     -

errors: No known data errors

pool: rpool
state: ONLINE
scan: none requested
config:

    NAME                STATE                READ WRITE CKSUM
    rpool                ONLINE                0     0     0
      c3t0d0s0           ONLINE                0     0     0

errors: No known data errors

NAME                USED  AVAIL  REFER  MOUNTPOINT
pool1                138K  7.81G   32K    /pool1
pool1/prod_data      31K   7.81G   31K    /prod_data
rpool               11.6G  4.04G  34.5K   /rpool
rpool/ROOT           9.94G  4.04G   31K    legacy
rpool/ROOT/solaris   9.94G  4.04G  9.70G   /
rpool/dump            630M   4.05G   611M   -
rpool/export         6.07M  4.04G   32K    /export
rpool/export/home    6.04M  4.04G   32K    /export/home
```

Identify the correct procedure for breaking the `/prod_data` mirror, removing `c4t0d0` and `c4t2d0`, and making the data on `c4t0d0` and `c4t2d0` accessible under the `dev_data` mount point.

- A. `zpool split pool1 pool2 c4t0d0 c4t2d0`  
`zpool import pool2zfs set mountpoint = /dev_data pool2/prod_data`
- B. `zpool detach pool1 pool2`  
`zpool attach pool2zfs 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 255**

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 260**

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 262**

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 267**

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 269

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 271

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 274

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

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

**Answer:** AF

**Explanation:** The Internet Control Message Protocol (ICMP) is one of the core protocols of the Internet Protocol Suite.

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

#### NEW QUESTION 276

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= YES` Restart the syslogd daemon.
- C. `svcadm enable svc:/system/system -log/config/log_from_remote` Restart the syslogd daemon.
- D. `svccfg -s svc:/system/system-log setprop config/log_from_remote=true` Restart the syslogd daemon.
- E. Set the following parameter in the `/etc/default/syslogd` file: `LOG_FROM_REMOTE=YES` Restart 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 278

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 280

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|lpr] [-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 285

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 289

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 292

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 297

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

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

**Answer:** ABC

**Explanation:** How to Monitor All Failed Login Attempts

This procedure captures in a syslog file all failed login attempts. Assume the Primary Administrator role, or become superuser.

1. Assume the Primary Administrator role, or become superuser.
2. (C) Set up the /etc/default/login file with the desired values for SYSLOG and SYSLOG\_FAILED\_LOGINS
3. (A) Create a file with the correct permissions to hold the logging information. Create the authlog file in the /var/adm directory.
4. (B) Edit the syslog.conf file to log failed password attempts.

### NEW QUESTION 301

You need to set up a local package repository to serve 75 client systems. Multiple clients will be using the package repository concurrently and you need to ensure that the local repository performs very well under this heavy load, especially during package intensive operations.

Which option would ensure the best performance of the repository during package-intensive operations by multiple clients?

- A. Set up multipathing on the package repository server to distribute the network load multiple network interfaces.
- B. Deploy a second instance of the package repository server to run as a read-writable mirror.
- C. Deploy a second instance of the package repository server to run as a read-only mirror.
- D. Deploy a second instance of the package repository server to run as a clone of the primary repository server.
- E. Deploy a package repository locally on each client.

**Answer:** A

### NEW QUESTION 303

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 308

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 310

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 312

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 313

User jack logs in to host solaris and then attempts to log in to host oracle using ssh. He receives the following error message:

The authenticity of host oracle (192.168.1.22) can't be established. RSA key fingerprint is 3B:23:a5:6d:ad:a5:76:83:9c:c3:c4:55:a5:18:98:a6

Are you sure you want to continue connecting (yes/no)?

Which two are true?



- A. The public host key supplied by solaris is not known to the host oracle.
- B. The error would not occur if the RSA key fingerprint shown in the error message was added to the /etc/ssh/known\_hosts file on solaris.
- C. The private host key supplied by oracle is not known to solaris.
- D. If jack answers yes, the RSA public key for the host oracle will be added to the known\_hosts file for the user jack.
- E. The public host key supplied by oracle is not known to the host solaris.

**Answer:** BD

**Explanation:** The fingerprints are used to guard against man in the middle attacks. Since ssh logins usually work over the internet (an insecure connection), someone could hijack your connection. When you try to log into yourmachine.com, he could get "in the middle" and return your challenge as if he was yourmachine.com. That way, he could get hold of your login password.

To make this attack harder, ssh stores the fingerprint of the server's public key on the first connection attempt. You will see a prompt like:

The authenticity of host 'eisen (137.43.366.64)' can't be established. RSA key fingerprint is cf:55:30:31:7f:f0:c4:a0:9a:02:1d:1c:41:cf:63:cf. Are you sure you want to continue connecting (yes/no)

When you enter yes, ssh will add the fingerprint to your known\_hosts file. you will see

Code:

Warning: Permanently added 'eisen, 137.43.366.64' (RSA) to the list of known hosts.

The next time you login, ssh will check whether the host key has changed. A changing host key usually indicates a man in the middle attack, and ssh refuses to connect.

#### NEW QUESTION 315

You have set up the task.max-lwps resource control on your Solaris 11 system.

Which option describes how to configure the system so that syslogd notifies you when the resources control threshold value for the task.max-lwps resource has been exceeded?

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

**Answer:** A

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

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

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

```
#
```

#### NEW QUESTION 316

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 319

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 321

You are asked to troubleshoot networking issues on an unfamiliar system. Select the correct command to display what network devices are installed.

- A. ifconfig -a

- B. dladm show-dev
- C. dladm show-phys
- D. dladm show-ether
- E. netadm show-dev
- F. netadm show-ether

**Answer:** C

#### NEW QUESTION 323

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

```
jack@solaris:/export/home/jack $ 1s 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 327

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 332

.....

## 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>