

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

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 5

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 id 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 6

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

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

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

Answer: B

Explanation: Test the remote connection with traceroute.

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

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

NEW QUESTION 7

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

```
cp bigfile /tmp
```

Will the command execute successfully?

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

Answer: A

Explanation: UFS quotas enable system administrators to control the size of file systems. Quotas limit the amount of disk space and the number of inodes, which roughly corresponds to the number of files, that individual users can acquire. For this reason, quotas are especially useful on the file systems where user home directories reside. As a rule, the public and

/tmp file systems usually do not benefit significantly by establishing quotas. Note: The cp command copies files and directories.

NEW QUESTION 8

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 9

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

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

NEW QUESTION 10

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

Which option correctly represents the output?

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

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

Answer: B

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

* who am i

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

user1 console Dec 30 20:20

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

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

Note:

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

Real UID

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

Effective UID

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

Saved UID

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

NEW QUESTION 10

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

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

Answer: A

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

crontab -e

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

NEW QUESTION 14

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 15

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 17

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 22

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 26

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 27

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

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

Answer: CE

NEW QUESTION 29

View the following information for a software package:

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

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

- A. pkg list gzip
- B. pkg info -r gzip
- C. pkg search -1 gzip
- D. pkg verify -v gzip
- E. pkg contents gzip

Answer: B

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

Example:

Listing Information About an Uninstalled Package

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

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

```
Version: 5.2.17
```

```
Build Release: 5.11
```

```
Branch: 0.175.0.0.0.1.530
```

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

```
Size: 44.47 MB
```

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

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

installed packages in the image.

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

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

NEW QUESTION 31

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 35

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 39

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

To accomplish this, you will create .

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

Answer: D

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

NEW QUESTION 44

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 47

On which is the open boot prom available?

- A. x86 only
- B. x86 64-Bit only
- C. SPARC only
- D. both x86 and x86 64-Bit
- E. x86, x86 64-Bit and SPARC

Answer: C

Explanation: No OpenBoot Environment on the Intel Platform. The Intel environment has no OpenBoot PROM or NVRAM. On Intel systems, before the kernel is started, the system is controlled by the basic input/output system (BIOS), the firmware interface on a PC. Therefore, many features provided by OpenBoot are not available on Intel systems.

Note: The Open Boot PROM (OBP) bootloader only exists within SPARC. Before Solaris 10 01/06, the bootloader for Solaris x86 was a Sun customized bootstrap software. After Solaris 10 01/06, it uses GRUB, a well known bootloader that's commonly used in the Linux world.

With GRUB, it's much easier to make the system dual-boot Linux and Solaris. GRUB extends the capabilities of the bootloader that was not available previously such as the ability to boot from a USB DVD drive. Those who have used Linux will be quite familiar with GRUB and its options.

NEW QUESTION 50

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 54

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

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

Answer: C

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

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

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

NEW QUESTION 55

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

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

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

Answer: CE

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

Example:

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

```
LINK IPKTS RBYTES INTRS POLLS CH<10 CH10-50 CH>50 e1000g0 101.91K 32.86M 87.56K 14.35K 3.70K 205 5
```

```
nxge1 9.61M 14.47G 5.79M 3.82M 379.98K 85.66K 1.64K vnic1 8 336 0 0 0 0
e1000g0 0 0 0 0 0 0
nxge1 82.13K 123.69M 50.00K 32.13K 3.17K 724 24
vnic1 0 0 0 0 0 0
```

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

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

List of supported RX fields: link

usedby

ibytes ipkts intrs polls

hdrops: hardware drops

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

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

NEW QUESTION 58

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/R00T/solaris	-	-	26.16M	static	2011-11-28 11:09
rpool/R00T/solaris/var	-	-	69.0K	static	2011-11-28 13:49
rpool/R00T/solaris/var@2011-11-28-18:49:38	-	-	0	static	2011-11-28 14:09
rpool/R00T/solaris/var@2011-11-28-19:09:23	-	-	975.0K	static	2011-11-28 12:29
rpool/R00T/solaris/var@install	-	-	70.0K	static	2011-11-28 13:49
rpool/R00T/solaris@2011-11-28-18:49:38	-	-	0	static	2011-11-28 14:09
rpool/R00T/solaris@2011-11-28-19:09:23	-	-	929.5K	static	2011-11-28 12:29
rpool/R00T/solaris@install	-	-	2.0K	static	2011-11-28 13:49
solaris-1	!	-	1.0K	static	2011-11-28 13:49
rpool/R00T/solaris-1	-	-	57.0K	static	2011-11-28 14:09
rpool/R00T/solaris-1/var	-	-	1.0K	static	2011-11-28 14:09
z1BE	-	-	-	-	-
rpool/R00T/z1BE	-	-	-	-	-
rpool/R00T/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

NEW QUESTION 61

Which two accurately identify features of a Solaris 10 branded zone?

- A. executes in a Solaris 10 global zone
- B. is created by importing a Solaris 10 flash archive
- C. enables Linux binary applications to run unmodified
- D. provides a complete runtime environment for Solaris 9 applications
- E. allows a Solaris 10 global zone to be migrated into a Solaris 11 system

Answer: BE

Explanation: B: It can be created by importing a Solaris 10 flash archive.

You can use the Oracle Solaris Flash archiving tools to create an image of an installed system that can be migrated into a zone.

The system can be fully configured with all of the software that will be run in the zone before the image is created. This image is then used by the installer when the zone is installed.

Note: You can use alternate methods for creating the archive. The installer can accept the following archive formats:

* cpio archives

* gzip compressed cpio archives

* bzip2 compressed cpio archives

* pax archives created with the -x xustar (XUSTAR) format ↑ ufsdump level zero (full) backups

Note:

Branded zones that run an environment different than the OS release on the system

* The lx branded zone introduced in the Solaris 10 8/07 release provides a Linux environment for your applications and runs on x86 and x64 machines on the Oracle Solaris 10 OS.

* The solaris8 and solaris9 branded zones enable you to migrate an Oracle Solaris 8 or Oracle Solaris 9 system to an Oracle Solaris 8 or Oracle Solaris 9 Container on a host running the Oracle Solaris 10 8/07 Operating System or later Oracle Solaris 10 release.

* The Oracle Solaris 10 Container brand is available in OpenSolaris build 127. These branded zones host Oracle Solaris 10 user environments.

Note: One of the powerful features of Solaris 11 is the ability to run a Solaris 10 environment in a zone. Solaris 10 allows you to run Solaris 8 and 9 environments in zones, but only on SPARC.

NEW QUESTION 65

Which operation will fail if the DNS configuration is incorrect?

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

Answer: E

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

NEW QUESTION 67

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 71

Which two options are characteristics of a fast reboot?

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

Answer: CD

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

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

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

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

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

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

The Fast Reboot feature is managed by SMF and implemented through a boot configuration service, svc:/system/boot-config. The boot-config service provides a means for setting or changing the default boot configuration parameters. When the config/fastreboot_default property is set to true, the system performs a fast reboot automatically, without the need to use the reboot -f command. This property's value is set to true on the x86 platform. For task-related information, including how to change the default behavior of Fast Reboot on the SPARC platform, see Accelerating the Reboot Process on an x86 Based System.

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

NEW QUESTION 73

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

- A. pkg verify p7zip
- B. pkg fix p7zip
- C. pkg rebuild-index p7zip
- D. pkg revert p7zip

- E. pkg uninstll p7zip
- F. pkg install p7zip
- G. pkg install --no-backup-be p7sip
- H. pkg refresh p7zip

Answer: D

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

NEW QUESTION 74

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 78

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 83

You have been asked to troubleshoot the initial configuration of a virtual network connecting two local zones with the outside world. View the exhibit.



The command
dladm create-vnic -1 vswitch192.168.1 vnic1 fails with the error
dladm: invalid link name 'vswitch192.168.1' What is the reason for this error?

- A. The name vswitch192.168.1 is not legal.
- B. The zone must be specified withdladm create-vnic -z zone3 vnic1.
- C. The virtual interface must be specified withdladm create-vnic -z zone3 vnic1.
- D. The virtual interface must be created withipadm create-vnic -1 switch192.168.1.
- E. The virtual switch must be created first withdladm create -etherstub vswitch192.168.1.

Answer: E

Explanation: There is no data-link named vswitch192.168. We need to create an etherstub first.

See Note and example below for details.

Note: Create a VNIC in the system's global zone.

```
# dladm create-vnic -l data-link vnic-name
```

data-link is the name of the interface where the VNIC is to be configured.

```
-l link, --link=link
```

link can be a physical link or an etherstub.

vnic-name is the name that you want to give the VNIC.

For example, to create a VNIC named vnic0 on interface e1000g0, you would type the following:

```
# dladm create-vnic -l e1000g0 vnic0
```

Example: Creating a Virtual Network Without a Physical NIC First, create an etherstub with name stub1:

```
# dladm create-etherstub stub1
```

Create two VNICs with names hello0 and test1 on the etherstub. This operation implicitly creates a virtual switch connecting hello0 and test1.

```
# dladm create-vnic -l stub1 hello0
```

```
# dladm create-vnic -l stub1 test1
```

NEW QUESTION 86

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 88

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 92

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

- A. Resources listed in /etc/dfs/dfstab are automatically shared on boot up.
- B. A directory cannot be shared if a subdirectory below it is already shared.
- C. Renaming a share created with the zfs set share command is not supported.
- D. NFS and SMB protocols cannot be used simultaneously to share the same directory.

Answer: AC

Explanation: A: ZFS can automatically share file systems by setting the sharenfs property. Using this property, you do not have to modify the /etc/dfs/dfstab file when a new file system is shared. The sharenfs property is a comma-separated list of options to pass to the share command. The value on is an alias for the default share options, which provides read/write permissions to anyone. The value off indicates that the file system is not managed by ZFS and can be shared through traditional means, such as the /etc/dfs/dfstab file. All file systems whose sharenfs property is not off are shared during boot.

NEW QUESTION 95

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

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

Answer: B

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

```
$ suPassword:
# crontab -l
#ident "@(#)root 1.19 98/07/06 SMI" /* SVr4.0 1.1.3.1 */
#
# The root crontab should be used to perform accounting data collection.
#
#
10 3 * * * /usr/sbin/logadm
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
30 3 * * * [ -x /usr/lib/gss/gsscred_clean ] && /usr/lib/gss/gsscred_clean
#10 3 * * * /usr/lib/krb5/kprop_script slave_kdcs
```

NEW QUESTION 100

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

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

Answer: ABD

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

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

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

NEW QUESTION 103

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 106

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 110

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

- A. While configuring the zone, add this entry: add rct1set name = capped.cpu-sharesadd value (priv = privileged, limit = 20, action = none)endexit
- B. While configuring the zone, add this entry: add rct1set name= zone.cpu-sharesadd value (priv=privileged, limit=20, action=none)endexitfrom command line, enter: # dispadmin- d FSS
- C. From the command line enter: #prct1 -n zone.cpu-shares - r - v 20 - i zone testzone
- D. From the command line, enter:#prct1 - n zone.cpu-shares - v 80 - r - i zone global

Answer: C

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

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

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

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

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

Note: project.cpu-shares

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

NEW QUESTION 111

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

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

A. /etc/security/policy.conf /etc/passwd

B. /etc/passwd /etc/shadow

C. /etc/security /etc/passwd

D. /etc/shadow /etc/passwd

Answer: B

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

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

NEW QUESTION 114

Examine this command and its output:

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

```
Name USED AVAIL REFER MOUNTPOINT
```

```
tank 2.41G 2.43G 32K /tank
```

```
tank/database 2.41G 2.43G 2.41G /tank/database tank/[email protected] 20K - 2.00G -
```

Next you execute:

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

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

A. It destroys the tank/database dataset.

B. It destroys tank/database and all descendant datasets.

C. It fails because the tank/ snapshot depends on the tank/database dataset.

D. It fails because the tank/ clone depends on the tank/database dataset.

E. It fails because the tank/database data set is not empty.

Answer: C

NEW QUESTION 116

The core dump configuration for your system is:

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

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

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

The server host name is: zeus

What will the per-process core file be named?

A. core.bash.2663.global

B. core.bash.2663.zeus

C. /var/core/core.bash.2663

D. /var/core/core.bash.2663.global

Answer: C

Explanation: Note the first line:

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

The program name is bash The runtime process ID is 2663

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

NEW QUESTION 121

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

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

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

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

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

Answer: A

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

How to Install the solaris10 Branded Zone

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

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

NEW QUESTION 122

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 123

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 124

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
K3eropfqsmPPY6uilyh9vpiFX2r1LHcbx139+vG6HOtVvuS8+lfMDtawm3WQvRuOopz3vVy
5GtMwtaOgehsXoT930Ryev1bH5myPtWKlipITsOd2sX9k3tvjrmme4KCGGss=
```

NEW QUESTION 127

Which option would you choose to display the kernel revision level for your operating system?

- A. ca
- B. /etc/release
- C. uname -a
- D. pkg info kernel
- E. banner (issued from the OpenBoot Prom)
- F. cat /etc/motd

Answer: B

NEW QUESTION 132

The global zone has 8 CPUS. YOU suspect that one of your non global /ones, dbzone, is consuming all of the CPU resources. Which command would you use to view the CPU utilization for all of the zones to confirm this?

- A. Run from the global zone:prstat -Z
- B. Run from each zone:login <zonename> mpstat
- C. Run from the global zone:zonestar -r summary
- D. Run from the global zone:rctladm -1
- E. Run from the global zone:prctl -i

Answer: A

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

NEW QUESTION 137

Identify the correct description of an IPS image.

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

Answer: D

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

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

NEW QUESTION 139

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

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

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

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

Answer: AC

Explanation: A: %PART

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

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

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

NEW QUESTION 144

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 145

In a fresh installation of Oracle Solaris 11, default datalinks are named with a generic 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 146

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 150

You are asked to determine user jack's default login directory. Which command would provide you with useful information?

- A. `cat /etc/passwd | grep jack`
- B. `cat /etc/group | grep jack`
- C. `cat /etc/shadow | grep jack`
- D. `cat /etc/default/passwd | grep jack`

Answer: A

Explanation: The `/etc/passwd` contains one entry per line for each user (or user account) of the system. All fields are separated by a colon (:) symbol. Total seven fields as follows.

1. Username: It is used when user logs in. It should be between 1 and 32 characters in length.
2. Password: An x character indicates that encrypted password is stored in `/etc/shadow` file.
3. User ID (UID): Each user must be assigned a user ID (UID). UID 0 (zero) is reserved for root and UIDs 1-99 are reserved for other predefined accounts. Further UID 100-999 are reserved by system for administrative and system accounts/groups.
4. Group ID (GID): The primary group ID (stored in `/etc/group` file)
5. User ID Info: The comment field. It allow you to add extra information about the users such as user's full name, phone number etc. This field use by finger command.
6. Home directory: The absolute path to the directory the user will be in when they log in. If this directory does not exists then users directory becomes /
7. Command/shell: The absolute path of a command or shell (`/bin/bash`). Typically, this is a shell. Please note that it does not have to be a shell.

NEW QUESTION 155

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 160

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

You have these requirements:

1. Users must have access to the data during the migration, so you cannot shutdown the file system while the migration takes place.
 2. Because you want to copy the data as quickly as possible, you need to increase the server resources devoted to the ZFS migration.
- Which method would you use to modify the ZFS shadow migration daemon defaults to increase the concurrency and overall speed of migration?

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

Answer: A

Explanation: shadowd is a daemon that provides background worker threads to migrate data for a shadow migration. A shadow migration gradually moves data from a source file system into a new "shadow" file system. Users can access and change their data within the shadow file system while migration is occurring. The shadowd service is managed by the service management facility, `smf(5)`.

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

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

`config_params/shadow_threads`

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

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

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

NEW QUESTION 163

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 164

You are installing the Solaris 11 Operation System by using the Text Installer. A panel prompts you to create a root password and a user account. Which four describe your options for completing this panel of the Installation?

- A. Creating a user account is optional.
- B. The root password must be set and cannot be blank.
- C. The root password can be left blank.
- D. If you provide a username, that user is assigned the root role.
- E. If you provide a username, that user is given root privileges.
- F. If you provide a username, root is an account rather than a role and is set to expire immediately.
- G. If you do not provide a username, root is an account rather than a role and is set to expire immediately.

Answer: ABDG

Explanation: A: You are not required to create a user account. B: You must create a root password. D: If you create a user account in this panel, you need to provide both the user's password and a root password. In this case, root will be a role assigned to the user. G: If you do not create a user account, you still need to provide a root password. In this case, root will be a regular user.

NEW QUESTION 166

The interface net3 should be operating, but is not. Command:

```
ipadm show-addr | grep net3
Response:
net3/v4          static      down      192.168.0.200/24
```

Which command should you enter next?

- A. ipadm create-ip
- B. ipadm enable-if
- C. ipadm show-if
- D. ipadm up-addr

Answer: B

Explanation: Enable-if -t interface
Enables the given interface by reading the configuration from the persistent store. All the persistent interface properties, if any, are applied and all the persistent addresses, if any, on the given interface will be enabled.
-t, --temporary
Specifies that the enable is temporary and changes apply only to the active configuration.

NEW QUESTION 170

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 175

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

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

Answer: C

Explanation: REPOSITORY DESCRIPTION

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

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

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

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

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

NEW QUESTION 177

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

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

Answer: B

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

NEW QUESTION 180

Identify three differences between the shutdown and init commands.

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

Answer: ABE

NEW QUESTION 184

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 185

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 186

`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 190

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 194

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

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

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

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

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

```
jack@solaris: useradd
```

What will happen, and why?

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

Answer: D

NEW QUESTION 196

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 200

A datalink can best be described as .

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


Answer: D

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

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

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

Network Stack Showing Network Devices, Links, and Interfaces:



NEW QUESTION 204

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

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

You then enter: `ipadm create-ip net3`

What is the output?

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

Answer: B

Explanation: According to the exhibit the interface already exists.
The command `ipadm create-ip net3` is supposed to create a new interface `net3`.

NEW QUESTION 208

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

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

Answer: BC

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

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

C: How to Install the solaris10 Branded Zone

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

NEW QUESTION 209

_____ 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 212

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

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

Answer: C

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

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

```
# ipadm show-addr
```

```
ADDROBJ TYPE STATE ADDR
```

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

NEW QUESTION 214

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

Which method allows you to meet both requirements?

1. Copy live data from `/production_ufs` to `/production_zfs` while `/production_ufs` is in use.
2. When the copy is complete, `/production_zfs` will contain an up-to date copy of `/production_ufs`

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

Answer: E

Explanation: Migrating Data With ZFS Shadow Migration

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

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

* A local or remote ZFS file system to a target ZFS file system

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

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

* Create an empty ZFS file system.

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

For example:

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

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

NEW QUESTION 218

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

```
export dennie_says
```

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

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

```
Echo $dennis_says
```

What will he see, and why?

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

Answer: A

Explanation: The `$HOME/.profile` file is an initialization file that is executed after the `/etc/profile` when logging in to the Bourne or Korn shell. The file contains user preferences for variable settings. If the ENV variable is set to `.kshrc`, the `.kshrc` file executes every time a new shell begins execution. The `$HOME/.profile` is copied from the `/etc/skel/local.profile` file by the Administration Tool when creating a new account.

Note: `/etc/skel/local.profile`

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

NEW QUESTION 220

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

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

Answer: BEF

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

slice as the install target, for example.

Note: An AI manifest provides installation instructions.

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

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

NEW QUESTION 223

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 227

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 232

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

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

Answer: BDFGI

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

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

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

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

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

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

```
-d imagepath
```

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

expanded at this location.

NEW QUESTION 237

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

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

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

Answer: B

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

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

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

In short,

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

NEW QUESTION 238

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

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

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

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

Answer: DF

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

Not E: According to line IGNORE dirmtime

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

```
CHECK all IGNORE dirmtime
```

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

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

NEW QUESTION 241

The current ZFS configuration on server is:

```
pool1/data@now          0      -      31K  -
pool1/data@monday      0      -      31K  -
rpool/ROOT/solaris@install 280M   -      3.40G -

pool1  3.97G   200K   3.97G   0%   1.00x   ONLINE  -
pool3  7.94G   6.25G   1.69G   78%   1.00x   ONLINE  -
rpool  15.9G   11.6G   4.24G   73%   1.00x   ONLINE  -

pool1          200K   3.91G   31K   /pool1
pool1/data     31K    3.91G   31K   /data
pool1/data2    18K    3.91G   31K   /data2
pool3          6.25G   1.56G   32K   /pool3
pool3/IPS      6.25G   1.56G   6.25G /pool3/IPS
```

You need to backup the /data file system while the file system is active.

Select the option that creates a full backup of the /data file system and stores the backup on server in the pool named backup.

- A. Mount -F nfs system: /backup / mntzfs snapshot pool/data@monday>/mnt/Monday
- B. Mount -F nfs systemB: /backup/mntzfs snapshot pool1/data@Mondayzfs clone pool1/data@monday/mnt/Monday
- C. Zfs send pool1/data@Monday | ssh system zfs rcv backup/monday
- D. Zfs snapshot pool1/data@Monday | ssh system zfs rcv backup/monday

Answer: C

Explanation: http://docs.oracle.com/cd/E23823_01/html/819-5461/ghzvz.html

NEW QUESTION 243

The default publisher on your system is:

```
PUBLISHER  TYPE      STATUS  URI
solaris    origin  online  http://pkg.oracle.com/solaris/release
```

You want to update the Oracle Solaris 11 environment on your system, but you are not able to connect this system to the Internet to access the default Oracle repository. A repository has been created on your local network and is named `http://server1.example.com`.

Which command would you choose to connect your system to the local repository?

- A. `pkg publisher` to specify the new publisher
- B. `pkg set-publisher` to set the stickiness on the `http://server1.example.com` publisher and unset stickiness for `http://pkg.oracle.com/solaris/release`
- C. `pkg add-publisher` to add the new publisher
- D. `pkg set-publisher` to set the origin for the publisher

Answer: D

Explanation: Solaris 11 Express makes it pretty easy to set up a local copy of the repository.

A common reason folks need access to a local repository is because their system is not connected to the Internet.

The `pkg set-publisher` command can be used to for example add a publisher or to enable or disable a publisher.

Note: Example Adding a Publisher

Use the `-g` option to specify the publisher origin URI.

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

Example Specifying the Preferred Publisher

Use the `-P` option to specify a publisher as the preferred publisher. The specified publisher moves to the top of the search order. You can specify the `-P` option when you add a publisher or you can modify an existing publisher.

```
# pkg set-publisher -P example.com
```

Example Enabling or Disabling a Publisher

Use the `-d` option to disable a publisher. The preferred publisher cannot be disabled. A disabled publisher is not used in package operations such as `list` and `install`.

You can modify the properties of a disabled publishers.

Use the `-e` option to enable a publisher.

```
# pkg set-publisher -d example2.com
```

NEW QUESTION 246

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

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

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

Answer: A

NEW QUESTION 251

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

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


Answer: AD

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

D: A location comprises certain elements of a network configuration, for example a name service and firewall settings, that are applied together, when required.

You can create multiple locations for various uses. For example, one location can be used when you are connected at the office by using the company intranet.

Another location can be used at home when you are connected to the public Internet by using a wireless access point. Locations can be activated manually or automatically, according to environmental conditions, such as the IP address that is obtained by a network connection.



NEW QUESTION 252

Which two SMF milestones can be specified at boot time?

- A. none
- B. network
- C. all
- D. config
- E. unconfig
- F. devices

Answer: AC

Explanation: The milestones that can be specified at boot time are none
single-user multi-user
multi-user-server all

NEW QUESTION 255

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

```
pgrep userprogram l5317
```

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

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

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

Answer: B

Explanation: You can use the `pgrep` and `pkill` commands to identify and stop command processes that you no longer want to run. These commands are useful when you mistakenly start a process that takes a long time to run.

To terminate a process:

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

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

NEW QUESTION 258

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 261

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 265

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 266

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

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

Answer: B

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

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

NEW QUESTION 267

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

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

Which option describes the solution to the problem?

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

Answer: C

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

STMF is controlled by stmfadm, and stmfadm is the majority of the commands you will be using to administer COMSTAR (COmmon Multiprotocol Scsi TARget). Install the packages you need for COMSTAR with iSCSI and reboot:

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

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

The COMSTAR target mode framework runs as the stmf service. By default, the service is disabled. You must enable the service to use COMSTAR functionality. You can identify the service with the svcs command. If you have not rebooted the server since installing the group/feature/storage-server package, the service might not be enabled correctly.

NEW QUESTION 270

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 274

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 275

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

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

Answer:

A

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

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

command returns the names of packages that match the query.

NEW QUESTION 276

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 277

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 279

.....

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>