

# Exam Questions 1Z0-821

Oracle Solaris 11 System Administrator

<https://www.2passeasy.com/dumps/1Z0-821/>



**NEW QUESTION 1**

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 2**

New features were added to ZFS in Oracle Solaris11. Your justification to upgrade from Solaris10 to Oracle Solaris11 is that it will be possible to take advantage of the enhancements that were made to ZFS.

Identify the three ZFS functions and features that are included in Oracle Solaris 11, but not in Solaris 10.

- A. Encrypted ZFS datasets
- B. Ability for ZFS to detect and remove redundant data from the tile system
- C. Shadow Data Migration
- D. Ability to split a mirrored ZFS storage pool
- E. Ability to use ZFS on the boot drive and boot to a ZFS root file system.
- F. elimination of the swap file system when using ZFS on the root disk

**Answer:** ABC

**Explanation:** <http://www.oracle.com/technetwork/server-storage/solaris11/overview/solaris-matrix-1549264.html>

**NEW QUESTION 3**

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

```

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

```

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

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

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

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

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

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

**Answer: C**

#### NEW QUESTION 4

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

```
NAME (PUBLISHER)VERSIONINFO
Compress/zip3.1.2-0.175.0.0.0.0.537if-
NAMEVERSIONDATECOMMENT
Compress/zip3.109 Dec 2011 04:50:38 ESTNone
```

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

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

**Answer: B**

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

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

#### NEW QUESTION 5

dbzone is currently running on your server.

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

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

**Answer: DE**

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

E: Use: `zlogin zone shutdown`

to cleanly shutdown the zone by running the shutdown scripts.

Use this procedure to cleanly shut down a zone.

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

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

**NEW QUESTION 6**

Review the zonestat command:

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

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

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

**Answer:** D

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

\* duration (here 24 h)

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

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

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

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

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

**NEW QUESTION 7**

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. Yes
- B. Quotas do not include any of the system files such as /tmp /swap.
- C. Yes
- D. The quota is set at the directory level, not the user level.
- E. No
- F. The command will fail because it will cause him to exceed his user quota.
- G. No
- 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**

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

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

**Answer:** CD

**NEW QUESTION 9**

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

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

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

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

Your next command should be:

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

**Answer: E**

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

```
# ipadm create-addr -T static -a address addrobj
```

where addrobj uses the naming format interface/user-defined-string, such as e1000g0/v4globalz.

Note:

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

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

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

#### NEW QUESTION 10

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

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

**Answer: D**

**Explanation:** pkg/image\_root

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

#### NEW QUESTION 10

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

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

**Answer: B**

**Explanation:** IPv6 defines three address types: unicast

Identifies an interface of an individual node.

multicast

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

anycast

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

#### NEW QUESTION 13

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

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

**Answer: BDE**

**Explanation:** B: By carrying SCSI commands over IP networks, the iSCSI protocol enables you to access block devices from across the network as if they were connected to the local system. COMSTAR provides an easier way to manage these iSCSI target devices.  
 D: Common Multiprotocol SCSI TARget, or COMSTAR, a software framework that enables you to convert any Oracle Solaris 11 host into a SCSI target device that can be accessed over a storage network by initiator hosts.  
 E: One IP port can handle multiple iSCSI target devices.

**NEW QUESTION 18**

On server A, you enter the following command to add a static route to serverA route -p add -host 192.168.1.101 192.168.1.101 -static  
 What is the purpose of this command?

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

**Answer:** B

**Explanation:** Note: # route -p add -host destination-IP gateway-IP -static where destination-IP and gateway-IP are IPv4 addresses of the host to be used as a target. For example, you would type the following to specify the target system 192.168.10.137, which is on the same subnet as the interfaces in IPMP group itops0:  
 \$ route -p add -host 192.168.10.137 192.168.10.137 -static  
 This new route will be automatically configured every time the system is restarted. If you want to define only a temporary route to a target system for probe-based failure detection, then do not use the -p option.

**NEW QUESTION 21**

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

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

**Answer:** AD

**Explanation:** A: the pkg list command display a list of packages in the current image, including state and other information. By default, package variants for a different architecture or zone type are excluded.  
 D: pkginfo displays information about software packages that are installed on the system (with the first synopsis, with -l) or that reside on a particular device or directory (with the second synopsis, with -r).  
 Without options, pkginfo lists the primary category, package instance, and the names of all completely installed and partially installed packages. It displays one line for each package selected.  
 With -r, retrieve the data from the repositories of the image's configured publishers. Note that you must specify one or more package patterns in this case.

**NEW QUESTION 22**

Review the boot environment information displayed on your system:

oldBE	-	-	149.0K	static	2011-11-28	15:15
newBE	!	-	363.05M	static	2011-11-28	14:47
solaris	-	-	100.68M	static	2011-11-20	18:09
solaris-1	NR	/	19.07G	static	2012-01-22	07:23

Which two options accurately describe the newBE boot environment?

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

**Answer:** BC

**Explanation:** If the boot environment is unbootable, it is marked with an exclamation point (!) in the Active column in the beadm list output.  
 The beadm command restricts actions on unbootable boot environments as follows: You cannot activate an unbootable boot environment. (B)  
 You cannot destroy a boot environment that is both unbootable and marked as active on reboot.  
 You cannot create a snapshot of an unbootable boot environment.  
 You cannot use an unbootable boot environment or boot environment snapshot with the -e option of beadm create.  
 You cannot rename an unbootable boot environment. (C)

**NEW QUESTION 24**

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

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

**Answer:** A

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

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

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

**NEW QUESTION 29**

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 33**

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 35**

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

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

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

**Answer:** CE

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

Example:

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

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

```
nxge1 9.61M 14.47G 5.79M 3.82M 379.98K 85.66K 1.64K vnic1 8 336 0 0 0 0
```

```
e1000g0 0 0 0 0 0 0
```

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

```
vnic1 0 0 0 0 0 0
```

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

-r  
 Display receive-side statistics only. Includes bytes and packets received, hardware and software drops, and so forth.  
 List of supported RX fields: link  
 iusedby  
 ibytes ipkts intrs polls  
 hdrops: hardware drops  
 sdrops: software drops (owing to bandwidth enforcement) ch<10: number of packet chains of length < 10  
 ch10-50: number of packet chains of length between 10 and 50 ch>50: number of packet chains of length > 50

**NEW QUESTION 38**

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	!R	-	2.0K	static	2011-11-28 13:49
solaris-1	-	-	1.0K	static	2011-11-28 13:49
rpool/R00T/solaris-1	-	-	-	-	-
rpool/R00T/solaris-1/var	-	-	57.0K	static	2011-11-28 14:09
z1BE	-	-	1.0K	static	2011-11-28 14:09
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](http://docs.oracle.com/cd/E23824_01/html/E21801/unbootable.html#scrolltoc)

**NEW QUESTION 39**

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 10 non-global zone on 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 43**

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 48

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 p7zip
- H. pkg refresh p7zip

**Answer:** D

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

#### NEW QUESTION 49

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

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

**Answer:** E

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

#### NEW QUESTION 54

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

Which option would you choose?

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

**Answer:** D

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

```
# zfs set reservation=5G tank/home/bill
# zfs get reservation tank/home/bill NAME PROPERTY VALUE SOURCE
tank/home/bill reservation 5G local
```

#### NEW QUESTION 56

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

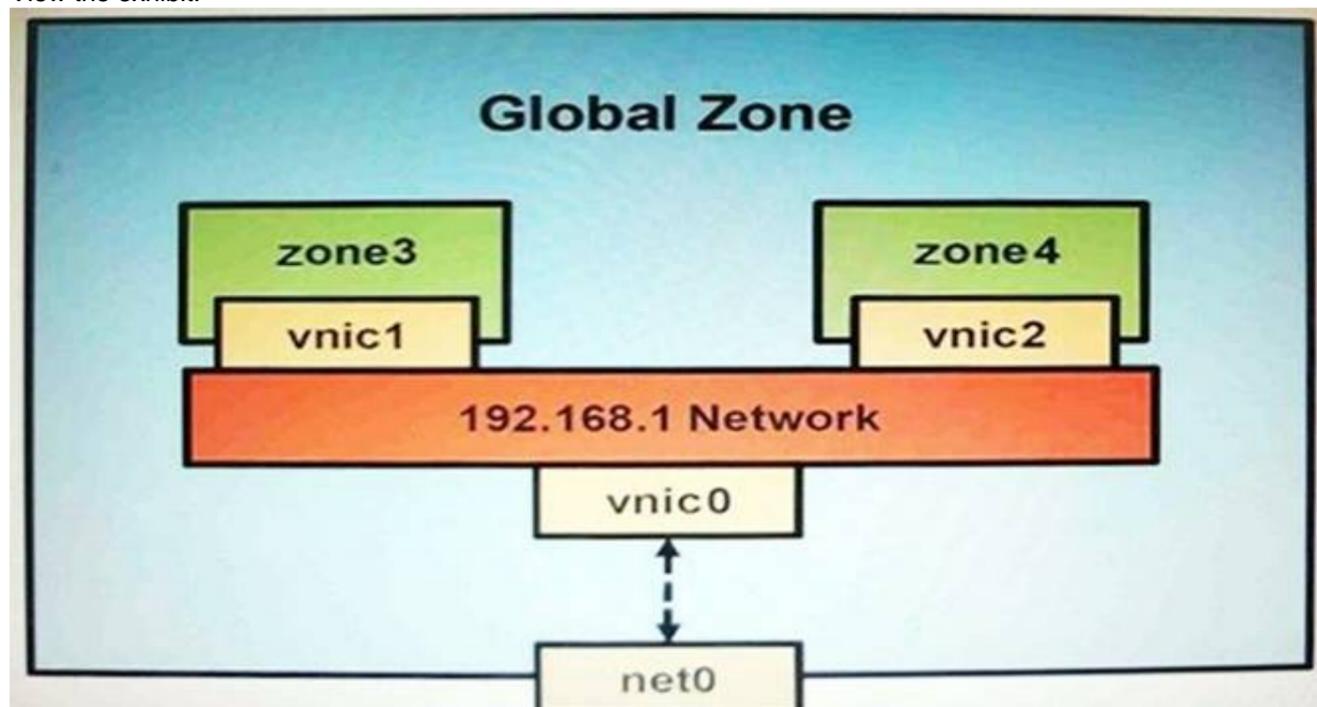
- A. It converts an Oracle Solaris 11 host into a SCSI target device that can be accessed over a storage network by Linux, Mac OS, or Windows client systems.
- B. iSCSI targets cannot be configured as dump devices.
- C. It provides support for iSCSI devices that use SLP.
- D. It is used to connect to 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 58**

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 -l 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 with `dladm create-vnic -z zone3 vnic1`.
- C. The virtual interface must be specified with `dladm create-vnic -z zone3 vnic1`.
- D. The virtual interface must be created with `ipadm create-vnic -l switch192.168.1`.
- E. The virtual switch must be created first with `dladm create -etherstub vswitch192.168.1`.

**Answer:** E

**Explanation:** There is no data-link named `vswitch192.168.1`. 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 60**

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 64

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 68

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 73

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 78

The following image properties are displayed on your system:

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

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

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

**Answer:** DF

**Explanation:** Image properties described below.

\* be-policy

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

/ default

Apply the default BE creation policy: create-backup.

/ always-new (D, F)

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

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

#### NEW QUESTION 79

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 83

Examine this command and its output:

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

```
Name USED AVAIL REFER MOUNTPOINT
```

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

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

Next you execute:

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

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

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

**Answer:** C

#### NEW QUESTION 88

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

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

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

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

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

**Answer:** A

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

How to Install the solaris10 Branded Zone

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

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

#### NEW QUESTION 89

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 94

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

- A. Run from the global zone:prstat -Z
- B. Run from each zone:zlogin <zonename> mpstat
- C. Run from the global zone:zonestar -r summary
- D. Run from the global zone:rctladm -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 95

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 96

You are using AI to install a new system. You have added to following information to the AI manifest:

```
<configuration type= "zone" name= "dbzone"
source = "http://sysA.example.com/zone_cfg/zone.cfg"/>
```

 Which statement is true with regard to the zone.cfg?

- A. The zone.cfg file is text file in a zonecfg export format.
- B. The zone.cfg file is an AI manifest that specifies how the zone is to be installed.
- C. The zone.cfg file is an xml file in a form suitable for use as a command file for the zonecfg command.
- D. The zone.cfg file is an SC profile with keywords that are specific for configuring a as part of the installation process.
- E. It is an xml configuration file from the /etc/zone director
- F. It will be used as a profile for the zon
- G. It specifies the zonename, zonepath, and other zonecfg parameters.

**Answer:** A

**Explanation:** [https://docs.oracle.com/cd/E23824\\_01/html/E21798/glitd.html#scrolltoc](https://docs.oracle.com/cd/E23824_01/html/E21798/glitd.html#scrolltoc) [http://docs.oracle.com/cd/E23824\\_01/html/E21798/glitd.html#aizoneconf](http://docs.oracle.com/cd/E23824_01/html/E21798/glitd.html#aizoneconf)

#### NEW QUESTION 98

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 102

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 106

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

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

**Answer:** AD

#### NEW QUESTION 107

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

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

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

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

**Answer:** BF

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

`Route add default <ipaddress>` Example:

`Route add default 192.168.1.1`

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

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

`/etc/defaultrouter` Example:

`Echo 192.168.1.1 > /etc/defaultrouter`

#### NEW QUESTION 111

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 115

Which two capabilities are provided by the OpenBoot PROM?

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

**Answer:** BC

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

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

#### NEW QUESTION 116

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 117

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 119

You want to delete the IPv4 address on the interface net3. Which command should you use?

- A. ipadm delete-ip net3/v4
- B. ipadm down-addr net3/v4
- C. ipadm disable-if net3/v4
- D. ipadm delete-vni net3/v4
- E. ipadm delete-addr net3/v4
- F. ipadm deiete-ipv4 ner3/v4

**Answer:** E

**Explanation:** The ipadm delete-addr subcommand removes addresses from interfaces. To remove an address from the IPMP group, type the following command:

```
# ipadm delete-addr addrobj
```

The addrobj uses the naming convention inder-interface/user-string.

#### NEW QUESTION 123

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 126

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 128

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

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

**Answer:** C

**Explanation:** Single quotes are most strict. They prevent even variable expansion. Double quotes prevent wildcard expansion but allow variable expansion. For example:

```
#!/bin/sh echo $SHELL
echo "$SHELL"
echo '$SHELL' This will print:
/usr/bin/bash
/usr/bin/bash
$SHELL
```

#### NEW QUESTION 130

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 132

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

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

**Answer:** ABC

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

#### NEW QUESTION 136

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 140

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 144

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

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

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

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

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

```
jack@solaris: useradd
```

What will happen, and why?

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

**Answer: D**

#### NEW QUESTION 146

Examine this command and its output:

```
$ zfs list -r -t all tank  
Name USED AVAIL REFER MOUNTPOINT  
tank 3.00G 1.84G 32K /tank  
tank/database 3.00G 1.84G 2.00G /tank/database tank/[email protected] 1.00G - 2.00G -  
Which two conclusions can be drawn based on this output?
```

- A. The tank dataset consumes 3 GB of storage.
- B. The tank/[email protected] dataset consumes 1 GB of storage that is shared with its parent.
- C. The tank/[email protected] dataset consumes 1 GB of storage that is not shared with its parent.
- D. The tank/[email protected] dataset consumes 2 GB of storage that is shared with its child.
- E. The tank/[email protected] dataset consumes 2 GB of storage that is not shared with its child.

**Answer: AB**

#### NEW QUESTION 147

You are logged in as root to a newly installed Solaris 11 system. You issue the command useradd -d, and then examine the /usr/sadm/defadduser file. This file

includes the entry defshell=/bin/sh. Which shell will now be the default for the next account created?

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

**Answer:** A

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

- \* Shell changes - The default shell, /bin/sh, is now linked to ksh93. The default user shell is the Bourne-again (bash) shell.
- \* The legacy Bourne shell is available as /usr/sunos/bin/sh.
- \* The legacy ksh88 is available as /usr/sunos/bin/ksh from the shell/ksh88 package.
- \* Korn shell compatibility information is available in /usr/share/doc/ksh/COMPATIBILITY.

#### NEW QUESTION 150

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

Which method allows you to meet both requirements?

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

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

**Answer:** E

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

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

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

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

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

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

For example:

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

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

#### NEW QUESTION 153

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

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

**Answer:** C

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

```
# pkg update
```

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

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

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

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

- \* The default installation repository at [pkg.oracle.com/solaris/release](http://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 155

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

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

**Answer:** A

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

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

#### NEW QUESTION 156

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 imag
- 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 159

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

```

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

```

```

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

```

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

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

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

Answer: C

#### NEW QUESTION 161

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 162**

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 166**

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 recv backup/monday
- D. Zfs snapshot pool1/data@Monday | ssh system zfs recv backup/monday

**Answer:** C

**Explanation:** [http://docs.oracle.com/cd/E23823\\_01/html/819-5461/ghzvz.html](http://docs.oracle.com/cd/E23823_01/html/819-5461/ghzvz.html)

**NEW QUESTION 169**

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

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

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

B. `cat so1.repo.iso-a so1.repo.iso-b > /export/IPS/repo` Set the `pkg/inst_root` property to true and the `pkg/readonly` property to /export/IPSSet the preferred publisher by using `pkg set-publisher -G http://serverA.example.com/ \-g http://pkg/oracle.com/solaris/release/solaris`

C. `cat so1.repo.iso-a so1.repo.iso-b > so1.full.iso` Mount the ISO image and use the `rsync` command to extract the contents of the ISO file to /export/IPS/repoSet the `pkg/inst_root` property to /export/IPS/repo and the `pkg/readonly` property to trueSet the preferred publisher by using `pkg set-publisher solaris \-g http://pkg.oracle.com/`

D. `cat so1.repo, iso-a so1.repo.iso-b > /export/IPS/repo.iso` Mount the ISO image and copy the repo directory from the ISO image to /export/IPS/repoSet the `pkg/inst_root` property and the `pkg/readonly` property to /export/IPS/repoSet the preferred `pkg/inst_root` property by using `pkg set-publisher - G http://serverA.example.com/ \- g http://pkg.oracle.com/solaris.com/release/- p solaris`

**Answer:** A

#### NEW QUESTION 171

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

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

**Answer:** AC

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

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

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

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

#### NEW QUESTION 175

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 180

You need to update an OS image on a client. The `pkg publishers` command displays the wrong publisher with the wrong update:

`PUBLISHERTYPESTATUSURI`

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

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

Solaris originonline `http://sysA.example.com`

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

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

**Answer:** C

**Explanation:** You can use the `pkg set-publisher` command to change a publisher URI. Changing a Publisher Origin URI To change the origin URI for a publisher, add the new URI and remove the old URI. Use the `-g` option to add a new origin URI. Use the `-G` option to remove the old origin URI.

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

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

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

#### NEW QUESTION 181

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 184

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 188

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 193

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 197

You need to set up an Oracle Solaris 11 host as an iSCSI target so that the host's disk can be accessed over a storage network. The disk device is c3t4d0. Which six options describe the steps that need to be taken on this host to enable an iSCSI target?

- A. Create a ZFS file system named iscsi/target.
- B. Create a zpool named iscsi with disk device c3t4d0
- C. Create zfs volume named iscsi/target.
- D. Use the stmfadm command to create a LUN using /dev/zvol/rdisk/iscsi/target.
- E. Use the stmfadm command to create a LUN using iscsi/target.
- F. Use the stmfadm command to make the LUN viewable.
- G. Use the stmfadm command to make the volume viewable.
- H. Enable the svc:/network/iscsi/target:default Service.
- I. Use the itadm command to create the iSCSI target.

**Answer: BCD FHI**

**Explanation:** How to Create an iSCSI LUN

The following steps are completed on the system that is providing the storage device.

Example: target# zpool create sanpool mirror c2t3d0 c2t4d0 (C)2. Create a ZFS volume to be used as a SCSI LUN. (D)3. Create a LUN for the ZFS volume. Example:

```
target# stmfadm create-lu /dev/zvol/rdisk/sanpool/vol1
Logical unit created: 600144F0B5418B0000004DDAC7C10001
```

4. Confirm that the LUN has been created.

Example

```
target# stmfadm list-lu
LU Name: 600144F0B5418B0000004DDAC7C10001
```

(F) 5. Add the LUN view.

This command makes the LUN accessible to all systems.

```
target# stmfadm add-view 600144F0B5418B0000004DDAC7C10001 How to Create the iSCSI Target
```

This procedure assumes that you are logged in to the local system will contains the iSCSI target.

Note: The stmfadm command manages SCSI LUNs. Rather than setting a special iSCSI property on the ZFS volume, create the volume and use stmfadm to create the LUN.

(H) 1. Enable the iSCSI target service.

```
target# svcadm enable -r svc:/network/iscsi/target:default
```

(I) 2. Create the iSCSI target.

```
target# itadm create-target
```

#### NEW QUESTION 200

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

Which command is best?

- A. kill -2 1234
- B. kill -15 1234
- C. kill -9 1234
- D. kill -1 1234

**Answer: B**

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

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

#### NEW QUESTION 201

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

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

**Answer: C**

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

Example:

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

#### NEW QUESTION 202

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)

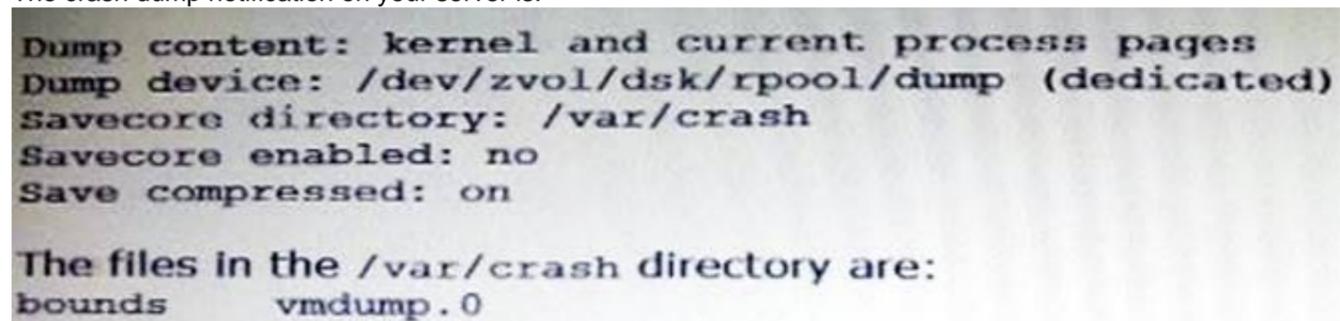
Plan 9 from Bell Labs 32

Minix 3 0 (Minix-vmd: 16)

QNX 6.4 8

#### NEW QUESTION 207

The crash dump notification on your server is:



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

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

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

Which command should you choose to display these two files?

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

**Answer:** A

**Explanation:** Decompress using `savecore -f vmdump.0`

`savecore -` save a crash dump of the operating system

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

#### NEW QUESTION 210

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

- A. `ppgsz`
- B. `pargs`

- C. pmap
- D. pgrep

**Answer:** B

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

**NEW QUESTION 212**

Which three options describe the purpose of the zonep2vchk command?

- A. Used on a Solaris 10 global zone to access the system for problems before migrating that system to a Solaris 10 branded zone.
- B. Used to access a Solaris 10 global zone for problems before migrating that zone to a Solaris 11 global zone
- C. Used to create zonecfg template for a Solaris 10 global zone that that will be migrated to a solaris10 branded zone.
- D. Used to migrate an Oracle Solaris 11 global zone to a non-global zone.
- E. Used to migrate a Solaris 10 global zone to a non-global zone on the same server; the non-global zone can then be migrated to a Solaris 11 server as a Solaris10 branded zone.

**Answer:** CDE

**Explanation:** zonep2vchk

- check a global zone's configuration for physical to virtual migration into non-global zone

The zonep2vchk utility is used to evaluate a global zone's configuration before the process of physical-to-virtual (p2v) migration into a non-global zone.

The p2v process involves archiving a global zone (source), and then installing a non-global zone (target) using that archive

Zonep2vchk serves two functions. First, it can be used to report issues on the source which might prevent a successful p2v migration. Second, it can output a template zonecfg, which can be used to assist in configuring the non-global zone target.

Zonep2vchk can be executed on a Solaris 10 or later global zone. To execute on Solaris 10, copy the zonep2vchk utility to the Solaris 10 source global zone.

When run on Solaris 10, a target release of S11 can be specified, which will check

for p2v into a Solaris 10 Branded zone.

**NEW QUESTION 217**

Review the non-global zone configuration displayed below:

```
zonename: dbzone
zonepath: /export/dbzone
brand: Solaris
autoboot: false
bootargs:
file-mac-profile:
pool:
limitpriv:
scheduling-class:
ip-type: exclusive
hostid:
fs-allowed:
anet:
    linkname: net0
    lower-link: auto
    allowed-address not specified
    configure-allowed-address: true
    defrouter not specified
    allowed-dhcp-cids not specified
    link-protection: mac-nospoof
    mac-address: random
    auto-mac-address: 2:8:20:97:40:20
    mac-prefix not specified
    mac-slot not specified
    vlan-id not specified
    priority nor specified
    rxrings not specified
    rxrings not specified
    mtu not specified
    maxlow not specified
    rxfanout not specified
```

The global zone has 1024 MB of physical memory. You need to limit the non-global zone so that it uses no more than 500 MB of the global zone's physical memory. Which option would you choose?

- A.  
From the global zone, issue these commands:  
zonecfg -z dbzone  
zonecfg:dbzone> set zone.max -memory=500m
- B.  
From the global zone, issue these commands:  
zonecfg -z dbzone  
zonecfg:dbzone>add rctl  
zonecfg:dbzone> set zone.max -memory=500m  
zonecfg:dbzone:capped-memory> end
- C.  
From the global zone, issue these commands:  
zonecfg -z dbzone  
zonecfg:dbzone> add capped-memory  
zonecfg:dbzone:capped-memory> set physical=500m  
zonecfg:dbzone:capped-memcry> end
- D.  
From the global zone, issue these commands:  
zonecfg -z dbzone  
zonecfg:dbzone> set physical=500m  
zonecfg:dbzone> end
- E.  
From the global zone, issue these commands:  
prctl -n zone.max -memory -v 500M -r -i dbzone

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

**Answer: C**

**Explanation:** Add a memory cap.

```
zonecfg:my-zone> add capped-memory
```

Set the memory cap.

```
zonecfg:my-zone:capped-memory> set physical=50m End the memory cap specification.
```

```
zonecfg:my-zone:capped-memory> end
```

#### NEW QUESTION 219

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

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

**Answer: ABDFH**

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

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

```
$ sudo pkg install slim_install
```

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

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

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

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

#### NEW QUESTION 221

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

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

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

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

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

**Answer: A**

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

Note: zpool attach [-f] pool device new\_device

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

### NEW QUESTION 222

Server A, Server B, and Server C are connected to the same network switch and are on the sari Each server has a single network interface, net0.

You received a tech support call that Server B has lost network connectivity. Your troubleshooting has discovered:

Server A can ping Server C, but not Server B. Server B can ping localhost, but not Server A or C. Server C can ping Server A, but not Server B.

On Server F3, you enter the following command: dladm show-phys | grep net0

Response:

```
net0/v4 Ethernet down 0 unknown el00gl
```

What is the next logical troubleshooting action?

- A. Run arp -a on all servers.
- B. Confirm that the router is working.
- C. Confirm that the power light of the network switch is on.
- D. Confirm that the physical network connections are intact.
- E. On Server A and C, run traceroute -n server.
- F. On Server B, run traceroute -n servera and traceroute -n serverc.

**Answer: D**

**Explanation:** Check the physical connection.

### NEW QUESTION 227

Consider the following commands:

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

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

- A. Hello, world
- B. cat: cannot open file1: No such file or directory Hello, world
- C. cat: cannot open file1: No such file or directory
- D. bash: syntax error near unexpected token '|'
- E. bash: syntax error broker pipe

**Answer: B**

### NEW QUESTION 230

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

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

Identify two correct statements.

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

**Answer:** AB

**Explanation:** Oracle Solaris provides predefined rights profiles. These profiles, listed in the /etc/security/prof\_attr, can be assigned by the root role to any account. The root role is assigned all privileges and all authorizations, so can perform all tasks, just as root can when root is a user.

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

```
$ su - root
```

```
Password: Type root password
```

```
#
```

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

#### NEW QUESTION 231

United States of America export laws include restrictions on cryptography.

Identify the two methods with which these restrictions are accommodated in the Oracle Solaris 11 Cryptographic Framework.

- A. Corporations must utilize signed X.509 v3 certificates.
- B. A third-party provider object must be signed with a certificate issued by Oracle.
- C. Loadable kernel software modules must register using the Cryptographic Framework SPI.
- D. Third-party providers must utilize X.509 v3 certificates signed by trusted Root Certification Authorities.
- E. Systems destined for embargoed countries utilize loadable kernel software modules that restrict encryption to 64 bit keys.

**Answer:** BC

**Explanation:** B: Binary Signatures for Third-Party Software

The elfsign command provides a means to sign providers to be used with the Oracle Solaris Cryptographic Framework. Typically, this command is run by the developer of a provider.

The elfsign command has subcommands to request a certificate from Sun and to sign binaries. Another subcommand verifies the signature. Unsigned binaries cannot be used by the Oracle Solaris Cryptographic Framework. To sign one or more providers requires the certificate from Sun and the private key that was used to request the certificate.

C: Export law in the United States requires that the use of open cryptographic interfaces be restricted. The Oracle Solaris Cryptographic Framework satisfies the current law by requiring that kernel cryptographic providers and PKCS #11 cryptographic providers be signed.

#### NEW QUESTION 235

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

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

You do not get an error message.

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

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

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

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

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

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

**Answer:** B

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

We need to add a password.

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

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

#### NEW QUESTION 240

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

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

Which three statements are correct?

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

**Answer:** ACD

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

#### NEW QUESTION 244

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

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

Which four correctly describe the output?

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

**Answer:** ACEG

**Explanation:** A: + Indicates the file/directory was added in the later dataset  
C: R Indicates the file/directory was renamed in the later dataset E: - Indicates the file/directory was removed in the later dataset  
G: M Indicates the file/directory was modified in the later dataset

Note: Identifying ZFS Snapshot Differences (zfs diff)

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

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

File or Directory Change Identifier

- \* File or directory is modified or file or directory link changed M
- \* File or directory is present in the older snapshot but not in the newer snapshot  
—
- \* File or directory is present in the newer snapshot but not in the older snapshot  
+
- \* File or directory is renamed R

#### NEW QUESTION 248

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

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

**Answer:** AF

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

#### NEW QUESTION 250

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

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

- A. svccfg -s svc:/system/system -log setprop start/exec= "syslogd -t"Restart the syslogd daemon.
- B. Set the following parameter in the /etc/syslogd.conf file: LOG\_FROM\_REMOTE= YESRestart the syslogd daemon.
- C. svcadm enable svc:/system/system -log/config/log\_from\_remoteRestart the syslogd daemon.
- D. svccfg -s svc:/system/system-log setprop config/log\_from\_remote=trueRestart the syslogd daemon.
- E. Set the following parameter in the /etc/default/syslogd file: LOG\_FROM\_REMOTE=YESRestart the syslogd daemon.

**Answer:** BD

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

LOG\_FROM\_REMOTE

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

#### NEW QUESTION 255

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

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

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

**Answer:** AEF

#### NEW QUESTION 258

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

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

These are the steps that you have followed:

1. Download and rsync the contents of the Oracle Solaris11 repository ISO image to the /repo directory.
2. Configure the repository server service properties. The svcprop command display, the IPS related properties:

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

The ls command displays the contents of the /repo directory:

```
#ls/repo
```

```
Pkg5.repository publisher
```

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

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

The pkg publisher command still displays: PUBLISHERTYPESTATUSURI

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

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

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

**Answer:** C

**Explanation:** Set the Publisher Origin To the File Repository URI

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

Example:

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

#### NEW QUESTION 262

The current ZFS configuration on your server is:

```
pool1 124K 3.91G 32K /pool1 pool1/data 31K 3.91G 31K /data
```

You need to create a new file system named /data2. /data2 will be a copy of the /data file system.

You need to conserve disk space on this server whenever possible.

Which option should you choose to create /data2, which will be a read writeable copy of the /data file system, while minimizing the amount of total disk space used in pool1?

- A. zfs set mountpoint=/data2 compression=on pool1/data2
- B. zfs snapshot pool1/data@nowzfs set mountpoint=/data2, compression=on pool1/data@now
- C. zfs create snapshot pool1/data@nowzfs send pool1/data@now | zfs recv pool1/data2
- D. zfs create snapshot pool1/data@nowzfs clone -o mountpoint=/data2 pool1/data@now pool1/data2
- E. zfs snapshot pool1/data@nowzfs clone -o mountpoint=/data2 -ocompression=on pool1/data@now pool1/data2
- F. zfs snapshot pool1/data@nowzfs clone -o mountpoint=/data2 pool1/data@now pool1/data2

**Answer:** E

**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

zfs clone [-p] [-o property=value] ... snapshot filesystem|volume Creates a clone of the given snapshot.

Note:

Because snapshots are fast and low overhead, they can be used extensively without great concern for system performance or disk use .

With ZFS you can not only create snapshot but create a clone of a snapshot.

A clone is a writable volume or file system whose initial contents are the same as the dataset from which it was created. As with snapshots, creating a clone is nearly instantaneous, and initially consumes no additional disk space. In addition, you can snapshot a clone.

A clone is a writable volume or file system whose initial contents are the same as the original dataset. As with snapshots, creating a clone is nearly instantaneous, and initially consumes no additional space.

Clones can only be created from a snapshot. When a snapshot is cloned, it creates an implicit dependency between the parent and child.

#### NEW QUESTION 267

View the Exhibit and review the file system information displayed from a remote server.

```

root@solaris:~# df -h
Filesystem      Size  Used  Available Capacity  Mounted on
rpool/ROOT/solaris 16G  3.4G   9.8G    26%      /
/devices         OK    OK    OK       0%      /devices
/dev             OK    OK    OK       0%      /dev
ctfs             OK    OK    OK       0%      /system/contract
proc            OK    OK    OK       0%      /proc
mnttab          OK    OK    OK       0%      /etc/mnttab
swap            1.1G  1.4M   1.1G     1%      /system/volatile
objfs           OK    OK    OK       0%      /system/object
sharefs         OK    OK    OK       0%      /etc/dfs/sharetab
/usr/lib/libc/libc_hwcapi.so.1
fd              13G  3.4G   9.8G    26%      /lib/libc.so.1
rpool/ROOT/solaris/var
swap            16G  183M   9.8G     2%      /var
swap            1.1G  48K   1.1G     1%      /tmp
rpool/export     16G  32K   9.8G     1%      /export
rpool/export/home 16G  32K   9.8G     1%      /export/home
rpool/export/home/bcalkins
rpool           16G  714K   9.8G     1%      /export/home/bcalkins
pool1/data      16G  39K   9.8G     1%      /rpool
pool1/data      1.0G  31K   1.0G     1%      /data
root@solaris:~#

```

You are configuring a new server. This new server has the following storage pool configured:

NAME	SIZE	ALLOC	FREE	CAP	DEDUP	HEALTH	ALTROOT
Pool1	15.9G	85K	15.9G	0%	1.00x	ONLINE	-

This new server also has the following file systems configured:

NAME	USED	AVAIL	REFER	MOUNTPOINT
pool1	85K	15.6G	31K	/pool1
rpool	5.81G	9.82G	39K	/rpool
rpool/ROOT	3.82G	9.82G	31K	legacy
rpool/ROOT/solaris	3.82G	9.82G	3.40G	/
rpool/ROOT/solaris/var	333M	9.82G	183M	/var
rpool/dump	970M	9.85G	940M	-
rpool/export	796K	9.82G	32K	/export
rpool/export/home	764K	9.82G	32K	/export/home
rpool/export/home/bcalkins	714K	9.82G	714K	/export/home/bcalkins
rpool/swap	1.03G	9.85G	1.00G	-

When you are finished building this new server, the pool1/data dataset must be an exact duplicate of note server. What is the correct procedure to create the pool1/data dataset on this new server?

- A. zfs create -o mountpoint=/data -o refquota=1g pool1/data
- B. zfs set mountpoint=none pool1zfs create pool1/data
- C. zfs set mountpoint=none pool1zfs create -o mountpoint=/data -o quota=1g pool1/data
- D. zfs create quota=1g pool1/data
- E. zfs create mountpoint=/data pool1/data
- F. zfs set quota=1g pool1/data

Answer: A

**NEW QUESTION 272**

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

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

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

Answer: BE

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

- \* username: jack
- \* encrypted password: \$5\$9JFrt54\$7JdwmO.F11Zt/ jFeeOhDmnw93LG7Gwd3Nd/cwCcNWFFg
- \* Last password change (lastchanged): Days since Jan 1, 1970 that password was last changed: 0
- \* Minimum: The minimum number of days required between password changes i.e. the number of days left before the user is allowed to change his/her password: 15
- Maximum: The maximum number of days the password is valid (after that user is forced to change his/her password): 30 Warn : The number of days before password is to expire that user is warned that his/her password must be changed: 3
- \* Inactive : The number of days after password expires that account is disabled
- \* Expire : days since Jan 1, 1970 that account is disabled i.e. an absolute date specifying when the login may no longer be used

**NEW QUESTION 276**

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

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

**Answer:** B

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

pkg info example:

```
$ pkg info p7zip Name: compress/p7zip
```

Summary: The p7zip compression and archiving utility

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

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

Build Release: 5.11

Branch: 0.175.0.0.0.2.537

Packaging Date: Wed Oct 19 09:13:22 2011

Size: 6.73 MB

FMRI: pkg://solaris/compress/p7zip@9.20.1, 5.11-0.175.0.0.0.2.537:20111019T091322Z

#### NEW QUESTION 279

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

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

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

**Answer:** A

#### NEW QUESTION 281

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

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

**Answer:** AEFG

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

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

form-<state>

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

<state>

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

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

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

Note 2:

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

#### NEW QUESTION 285

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

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

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

Which two are true?

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

**Answer:** BD

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

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

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

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

Code:

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

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

#### NEW QUESTION 287

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

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

Which option describes the state of this server?

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

**Answer:** C

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

#### NEW QUESTION 289

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

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

Which two statements are correct?

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

**Answer:** AC

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

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

#### NEW QUESTION 293

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

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

**Answer:** B

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

#### NEW QUESTION 297

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 1Z0-821 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 1Z0-821 Product From:

<https://www.2passeasy.com/dumps/1Z0-821/>

### Money Back Guarantee

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

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