

Exam Questions 1Z0-821

Oracle Solaris 11 System Administrator

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



NEW QUESTION 1

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

How can this be accomplished?

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

Answer: E

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

The value assigned by umask is subtracted from the default.

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

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

NEW QUESTION 2

View the Exhibit.

```

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

Which is true regarding the disk drive?

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

Answer: A

Explanation: Installing a ZFS Root Pool

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

NEW QUESTION 3

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 4

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 5

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 6

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 7

You attempted to reboot a system via the init command, however the system did not perform boot sequence into the Oracle Solaris Operating Environment. You are presented with a prompt from the OpenBoot PROM. Which command would you enter, to boot the system from the default device?

- A. boot -net install
- B. boot
- C. boot -default
- D. boot -s0

Answer: B

Explanation: Boot

With this form, boot loads and executes the program specified by the default boot arguments from the default boot device

Note: boot has the following general format: boot [device-specifier] [arguments]

where device-specifier and arguments are optional.

NEW QUESTION 8

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 9

You are going to use the Automated installer (AI) to install a non global zone named zone1. You have created a custom manifest for the non-global zone and named it zone1manifest

Which command will you use to add this custom manifest to the s11-sparc install service and associate this custom manifest with the non-global zone?

- A. installadm create-profile -n s11-sparc -f /tmp/zone1manifest.xml - c
- B. installadm create-manifest -n s11-sparc -f /tmp/zone1manifest.xml -m
- C. installadm create-client -n s11-sparc -f /tmp/zone1manifest.xml -m zone1manifest -c zonename= "zone1"
- D. installadm create-service - n s11-sparc -f /tmp/zone1manifest.xml -m zone1manifest - c zonename="zone1"

Answer: B

Explanation: installadm add-manifest

Associates manifests with a specific install service, thus making the manifests available on the network, independently from creating a service. When publishing a non-default manifest, it is required to associate criteria either via criteria entered on the command line (-c) or via a criteria XML file (-C).

NEW QUESTION 10

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

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

- A. `ipadm`
- B. `tracert`
- C. `dladm`
- D. `snoop`
- E. `arp`

Answer: B

Explanation: Test the remote connection with `tracert`.

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

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

NEW QUESTION 10

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

Select two unsupported items that will require changes.

- A. `terminal = zterms`
- B. `name_service=NTS+`
- C. `timezone=US/pacific`
- D. `system_locale=en_US`
- E. `network_interface=PRIMARY`
- F. `root_password=rJmvLUXM10cU`

Answer: AD

Explanation: A: `terminal`

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

D: `system_locale`

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

NEW QUESTION 11

`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. No
- 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 16

View the Exhibit.

```
<?xml version="1.0" encoding="UTF-8"?>
<!--
Copyright (c) 2011, Oracle and/or its affiliates. All rights reserved.
-->
<!DOCTYPE auto_install SYSTEM "file:///usr/share/install/ai.dtd.1">
<auto_install>
  <ai_instance name="zone_default">
    <target>
      <logical>
        <zpool name="rpool">
          <filesystem name="export" mountpoint="/export"/>
          <filesystem name="export/home"/>
          <be name="solaris">
            <options>
              <option name="compression" value="on"/>
            </options>
          </be>
        </zpool>
      </logical>
    </target>
    <software type="IPS">
      <destination>
        <image>
          <!-- Specify locales to install -->
          <facet set="false">facet.locale.*</facet>
          <facet set="true">facet.locale.de</facet>
          <facet set="true">facet.locale.de_DE</facet>
          <facet set="true">facet.locale.en</facet>
          <facet set="true">facet.locale.en_US</facet>
          <facet set="true">facet.locale.es</facet>
          <facet set="true">facet.locale.es_ES</facet>
          <facet set="true">facet.locale.fr</facet>
          <facet set="true">facet.locale.fr_FR</facet>
          <facet set="true">facet.locale.it</facet>
          <facet set="true">facet.locale.it_IT</facet>
          <facet set="true">facet.locale.ja</facet>
          <facet set="true">facet.locale.ja_*</facet>
          <facet set="true">facet.locale.ko</facet>
          <facet set="true">facet.locale.ko_*</facet>
          <facet set="true">facet.locale.pt</facet>
          <facet set="true">facet.locale.pt_BR</facet>
          <facet set="true">facet.locale.zh</facet>
          <facet set="true">facet.locale.zh_CN</facet>
          <facet set="true">facet.locale.zh_TW</facet>
        </image>
      </destination>
      <software_data action="install">
        <name>pkg:/group/system/solaris-small-server</name>
      </software_data>
    </software>
  </ai_instance>
</auto_install>
```

The file came from your Automated Installer (AI) install server. The file is .

- A. An AI SC profile for non-global zones
- B. The default AI config file for non-global zones
- C. The default AI manifest for non-global zones
- D. A custom AI manifest

Answer: D

Explanation: ai_manifest

- Automated installation manifest file format

Synopsis

/usr/share/install/ai.dtd.1

Some customizations have been made, such as the selection of specific locales.

NEW QUESTION 20

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

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

Answer: B

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

cat /etc/release

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

Example:

cat /etc/release

Oracle Solaris 10 8/11 s10s_u10wos_17b SPARC

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

NEW QUESTION 22

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 27

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

whoami;who am i;id

Which option correctly represents the output?

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

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

Answer: B

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

* who am i

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

user1 console Dec 30 20:20

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

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

Note:

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

Real UID

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

Effective UID

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

Saved UID

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

NEW QUESTION 28

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

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

Answer: A

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

crontab -e

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

NEW QUESTION 31

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 33

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

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

You then ran: ipadm show-if

The result indicated that the interface was down.

You then ran:

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

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

The last command indicated that the interface was up.

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

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

Answer: B

Explanation: The first IP address is already in use.

NEW QUESTION 38

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

```
PUBLISHERTYPESTATUSURI
```

```
solarisoriginonlinehttp://sysA.example.com
```

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

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

Answer: AE

Explanation: A: Configure the Repository Server Service

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

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

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

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

Restart the pkg.depotd repository service.

```
# svcadm refresh application/pkg/server
```

```
# svcadm enable application/pkg/server
```

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

NEW QUESTION 42

The advantage of core files is that they allow you an opportunity to examine the cause of problems, so that they can be resolved.

However, core files must be managed because they .

- A. take up large amounts of disk space
- B. make numerous entries into the /var/adm/wtmpx file
- C. steal resources from the processor, slowing down system performance
- D. fill up swap space; this will begin to slow the system due to swaps
- E. fill up swap space; this will begin to slow the system due to paging

Answer: A

Explanation: Part of the job of cleaning up heavily loaded file systems involves locating and removing files that have not been used recently. You can locate unused files by using the ls or find commands.
 Other ways to conserve disk space include emptying temporary directories such as the directories located in /var/tmp or /var/spool, and deleting core and crash dump files.
 Note: Core files are generated when a process or application terminates abnormally. Core files are managed with the coreadm command.
 For example, you can use the coreadm command to configure a system so that all process core files are placed in a single system directory. This means it is easier to track problems by examining the core files in a specific directory whenever a process or daemon terminates abnormally.

NEW QUESTION 46

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 47

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 50

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

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

Answer: CE

NEW QUESTION 55

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 60

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

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

Answer: D

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

NEW QUESTION 63

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 66

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 71

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 73

Which two accurately describe the Solaris IPS repository?

- A. It contains a collection of operating system patches.
- B. It contains a collection of software packages.
- C. All packages within an IPS package repository reside in a catalog.
- D. It is an ISO image of the Solaris installation media.
- E. The packages in a catalog are associated with a specific publisher.

Answer: BE

Explanation: Image Packaging System (IPS) is a new network based package management system included in Oracle Solaris 11. It provides a framework for complete software lifecycle management such as installation, upgrade and removal of software packages. IPS also enables you to create your own software packages, create and manage package repositories, and mirror existing package repositories.

Oracle Solaris software is distributed in IPS packages. IPS packages are stored in IPS package repositories, which are populated by IPS publishers.

E: The following command displays property information about the local repository.

```
$ pkgrepo get -s /export/repoSolaris11
```

```
SECTION PROPERTY VALUE publisher prefix solaris repository description This\ repository\ serves\ a\ copy\ of\ the\ Oracle\ Solaris\ 11\ Build\ 175b\ Package\
Repository. repository name Oracle\ Solaris\ 11\ Build\ 175b\ Package\ Repository
repository version 4
```

The value of the publisher prefix specifies that solaris is to be used in the following cases:

When more than one publisher's packages are present and no publisher is specified in the package name in the pkg command

When packages are published to the repository and no publisher is specified.

NEW QUESTION 78

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

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

Answer: C

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

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

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

NEW QUESTION 81

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

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

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

Answer: CE

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

Example:

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

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

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

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

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

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

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

-r

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

List of supported RX fields: link

usedby

ibytes ipkts intrs polls

hdrops: hardware drops

sdrops: software drops (owing to bandwidth enforcement) ch<10: number of packet chains of length < 10
 ch10-50: number of packet chains of length between 10 and 50 ch>50: number of packet chains of length > 50

NEW QUESTION 83

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	/	367.9/M	static	2011-11-28 11:09
rpool/ROOT/solaris	-	-	26.16M	static	2011-11-28 11:09
rpool/ROOT/solaris/var	-	-	69.0K	static	2011-11-28 13:49
rpool/ROOT/solaris/var@2011-11-28-18:49:38	-	-	0	static	2011-11-28 14:09
rpool/ROOT/solaris/var@2011-11-28-19:09:23	-	-	975.0K	static	2011-11-28 12:29
rpool/ROOT/solaris/var@install	-	-	70.0K	static	2011-11-28 13:49
rpool/ROOT/solaris@2011-11-28-18:49:38	-	-	0	static	2011-11-28 14:09
rpool/ROOT/solaris@2011-11-28-19:09:23	-	-	929.5K	static	2011-11-28 12:29
rpool/ROOT/solaris@install	!R	-	2.0K	static	2011-11-28 13:49
solaris-1	-	-	1.0K	static	2011-11-28 13:49
rpool/ROOT/solaris-1	-	-	-	-	-
rpool/ROOT/solaris-1/var	-	-	57.0K	static	2011-11-28 14:09
z1BE	-	-	1.0K	static	2011-11-28 14:09
rpool/ROOT/z1BE	-	-	-	-	-
rpool/ROOT/z1BE/var	-	-	-	-	-

Which two options describe the solaris-1 boot environment?

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

Answer: AE

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

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

N – Active now.

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

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

http://docs.oracle.com/cd/E23824_01/html/E21801/unbootable.html#scrolltoc

NEW QUESTION 87

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 92

Which operation will fail if the DNS configuration is incorrect?

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

Answer: E

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

NEW QUESTION 94

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

```
~/bigscript &
```

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

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

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

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

Answer: A

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

```
grep: id 15166
```

```
ps -aef 15165
```

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

NEW QUESTION 99

You have already generated a 256-bit AES raw key and named the keystore file /mykey. You need to use the key to create an encrypted file system.

Which command should you use to create a ZFS encrypted file system named pool1/encrypt using the /mykey keystore?

- A. zfs create -o encryption = /mykey pool1/encrypt
- B. zfs create -o encryption = 256-ccm -o keysource = raw, file : ///my key pool1/encrypt
- C. zfs create -o encryption = AES keysource = /mykey pool1/encrypt
- D. zfs create -o encryption = on keystore = /mykey pool1/encrypt

Answer: B

Explanation: Example: Encrypting a ZFS File System by Using a Raw Key

In the following example, an aes-256-ccm encryption key is generated by using the pktool command and is written to a file, /cindykey.file.

```
# pktool genkey keystore=file outkey=/cindykey.file keytype=aes keylen=256
```

Then, the /cindykey.file is specified when the tank/home/cindy file system is created.

```
# zfs create -o encryption=aes-256-ccm -o keysource=raw, file:///cindykey.file tank/home/cindys
```

NEW QUESTION 103

Consider the following commands:

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

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

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

Answer: B

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

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

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

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

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

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

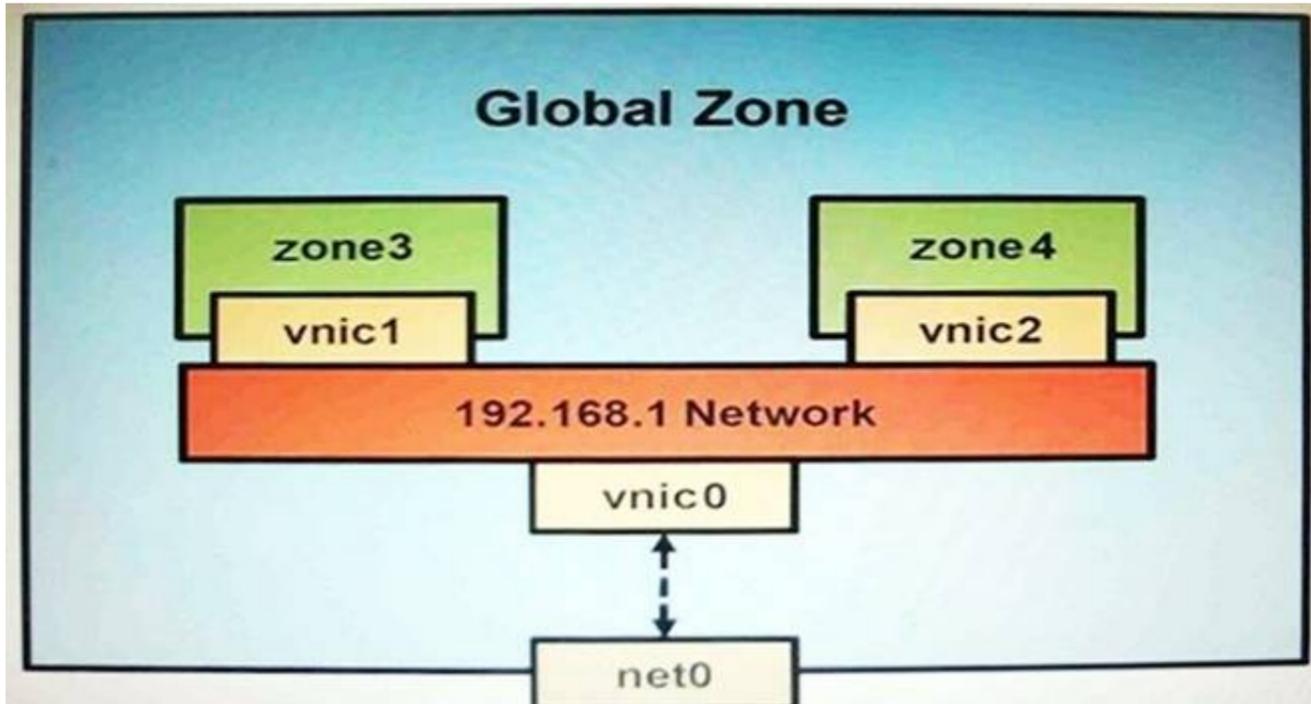
AND and OR lists are sequences of one or more pipelines separated by the control operators '&&' and '|', respectively. AND and OR lists are executed with left associativity.

An AND list has the form command1 && command2

command2 is executed if, and only if, command1 returns an exit status of zero.
 An OR list has the form command1 || command2
 command2 is executed if, and only if, command1 returns a non-zero exit status.
 The return status of AND and OR lists is the exit status of the last command executed in the list.
 Note #2 (on exit status): Zero means command executed successfully, if exit status returns non-zero value then your command failed to execute.

NEW QUESTION 108

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. 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 113

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 116

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 119

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

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

Answer: AC

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

NEW QUESTION 123

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 124

You want to configure an iSCSI target device on your system. Select the group package required to install this functionality on your system.

- A. `storage-server`
- B. `solaris-small-server`
- C. `storage-avs`
- D. `storage-nas`

Answer: A

Explanation: How to Create an iSCSI LUN

The disk volume provided by the server is referred to as the target. When the LUN is associated with an iSCSI target, it can be accessed by an iSCSI initiator. The following tasks are completed on the system that is providing the storage device.

1. Install the COMSTAR storage server software. `target# pkg install storage-server`
- Etc.

NEW QUESTION 128

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

- A. There are boot environments for zones.
- B. Administrators can delegate common administration tasks by using RBAC.

- C. Oracle Solaris 11 supports Solaris 8, 9, and 10 branded zones.
- D. You can migrate a physical Solaris 10 system and its non-global zones to a solaris10 branded zone running on an Oracle Solaris 11 system.
- E. It is possible to change the host ID of a zone.

Answer: ABD

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

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

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

NEW QUESTION 131

A change in your company's security policy now requires an audit trail of all administrators assuming the sysadm role, capturing: There are two command necessary to accomplish this change. One is a rolemod command. What is the other?

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

Answer: B

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

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

For all users and roles, add the AUE_PFEEXEC audit event to their preselection mask.

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

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

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

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

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

Note: [-t] -setpolicy [+|-]policy_flag[, policy_flag ...]

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

NEW QUESTION 132

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 134

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

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

Answer: C

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

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

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

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

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

Note: project.cpu-shares

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

NEW QUESTION 137

The core dump configuration for your system is:

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

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

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

The server host name is: zeus

What will the per-process core file be named?

- A. core.bash.2663.global
- B. core.bash.2663.zeus
- C. /var/core/core.bash.2663
- D. /var/core/core.bash.2663.global

Answer: C

Explanation: Note the first line:

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

The program name is bash The runtime process ID is 2663

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

NEW QUESTION 141

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. zone cfg - a s10-zone create - t SUNWsolaris10\</net/backup_server/s10-system.flar
- E. zonecfg -z s10-zone install -f /net/backup/backup_server/s10-system.flar

Answer: A

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

How to Install the solaris10 Branded Zone

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

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

NEW QUESTION 146

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

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

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

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

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

Answer: A

NEW QUESTION 150

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

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

Answer: C

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

NEW QUESTION 152

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 http://docs.oracle.com/cd/E23824_01/html/E21798/glitd.html#aizoneconf

NEW QUESTION 156

Which option displays the result of running the zfs list command?

- A)

```
NAME      SIZE  ALLOC   FREE   CAP  DEDUP  HEALTH  ALTROOT
pool1    15.9G 144K    15.9G   0%   1.00x  ONLINE  -
```
- B)

```
NAME      USED  AVAIL   REFER  MOUNTPOINT
pool1    144K  15.6G   31K    none
```
- C)

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

```
pool          capacity      operations      bandwidth
pool          alloc    free    read  write  read  write
-----
pool1         144K    15.9G     0     0     62    754
rpool        6.35G    9.52G     5     1   44.4K  10.6K
zone         3.41G    12.5G     0     0     76     17
-----
```

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

Answer: B

Explanation: The zfs list command provides an extensible mechanism for viewing and querying dataset information. You can list basic dataset information by using the zfs list command with no options. This command displays the names of all datasets on the system and the values of their used, available, referenced, and mountpoint properties. For more information about these properties, see *Introducing ZFS Properties*.

For example:

```
# zfs list
NAME USED AVAIL REFER MOUNTPOINT
pool 476K 16.5G 21K /pool
pool/clone 18K 16.5G 18K /pool/clone pool/home 296K 16.5G 19K /pool/home
pool/home/marks 277K 16.5G 277K /pool/home/marks pool/home/marks@snap 0 - 277K -
pool/test 18K 16.5G 18K /test
```

NEW QUESTION 157

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

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

Answer: B

Explanation: The User Datagram Protocol (UDP) is one of the core members of the Internet Protocol Suite, the set of network protocols used for the Internet. With UDP, computer applications can send messages, in this case referred to as datagrams, to other hosts on an Internet Protocol (IP) network without requiring prior communications to set up special transmission channels or data paths. Compared to TCP, UDP is a simpler message-based connectionless protocol.

NEW QUESTION 158

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 161

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 164

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 165

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 168

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 172

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

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

Answer: A

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

1. Username: It is used when user logs in. It should be between 1 and 32 characters in length.

2. Password: An x character indicates that encrypted password is stored in /etc/shadow file.

3. User ID (UID): Each user must be assigned a user ID (UID). UID 0 (zero) is reserved for root and UIDs 1-99 are reserved for other predefined accounts. Further UID 100-999 are reserved by system for administrative and system accounts/groups.

4. Group ID (GID): The primary group ID (stored in /etc/group file)

5. User ID Info: The comment field. It allow you to add extra information about the users such as user's full name, phone number etc. This field use by finger command.

6. Home directory: The absolute path to the directory the user will be in when they log in. If this directory does not exists then users directory becomes /

7. Command/shell: The absolute path of a command or shell (/bin/bash). Typically, this is a shell. Please note that it does not have to be a shell.

NEW QUESTION 173

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 174

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 175

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 176

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

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

Answer: C

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

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

`-g`

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

NEW QUESTION 179

Oracle Solaris 11 kernel encounters a fatal error, and it results in a system panic.

What type of file does this generate?

- A. `a.out`
- B. `objdump`
- C. core dump
- D. tape dump
- E. crash dump

Answer: C

Explanation: A kernel panic is a type of error that occurs when the core (kernel) of an operating system receives an instruction in an unexpected format or when it fails to handle properly. A kernel panic can also follow when the operating system can't recover from a different type of error. A kernel panic can be caused by damaged or incompatible software or, more rarely, damaged or incompatible hardware.

When a server kernel panics it abruptly halts all normal system operations. Usually, a kernel process named `panic()` outputs an error message to the console and stores debugging information in nonvolatile memory to be written to a crash log file upon restarting the computer. Saving the memory contents of the core and associated debugging information is called a “core dump.”

NEW QUESTION 180

Which three of the components could be used in a ZFS storage pool, but are not recommended configurations?

- A. A file on a UFS file system
- B. A Veritas Volume Manager (VxVM) volume
- C. A LUN in a hardware RAID array
- D. A disk slice from an SMI labeled disk
- E. A Solaris Volume Manager (SVM) volume
- F. An EFI labeled disk

Answer: ABE

Explanation: A: ZFS also allows you to use UFS files as virtual devices in your storage pool. This feature is aimed primarily at testing and enabling simple experimentation, not for production use. The reason is that any use of files relies on the underlying file system for consistency. If you create a ZFS pool backed by files on a UFS file system, then you are implicitly relying on UFS to guarantee correctness and synchronous semantics.

However, files can be quite useful when you are first trying out ZFS or experimenting with more complicated layouts when not enough physical devices are present. All files must be specified as complete paths and must be at least 64 Mbytes in size.

B, E: You can construct logical devices for ZFS using volumes presented by software-based volume managers, such as Solaris Volume Manager (SVM) or Veritas Volume Manager (VxVM). However, these configurations are not recommended. While ZFS functions properly on such devices, less-than-optimal performance might be the result.

NEW QUESTION 185

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

```
ipadm show-addr | grep net3
```

Response:

```
net3/v4          static      down      192.168.0.200/24
```

Which command should you enter next?

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

Answer: B

Explanation: Enable-if -t interface

Enables the given interface by reading the configuration from the persistent store. All the persistent interface properties, if any, are applied and all the persistent addresses, if any, on the given interface will be enabled.

-t, --temporary

Specifies that the enable is temporary and changes apply only to the active configuration.

NEW QUESTION 187

Review the information taken from your server:

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

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

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

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

Answer: A

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

-r

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

NEW QUESTION 191

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 195

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 196

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 197

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

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

Answer: BC

Explanation: B: banner

Displays power-on banner.

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

SPARCstation 2, Type 4 Keyboard

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

Displays version and date of the boot PROM.

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

NEW QUESTION 198

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 201

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 202

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

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

Answer: AD

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

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

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

NEW QUESTION 207

Which two are user definable OpenBoot parameters that can be set in the OpenBoot PROM?

- A. IP address for the system console
- B. Host ID
- C. System date and time
- D. Default boot device
- E. Verbose hardware diagnostics
- F. Powering off the hardware

Answer: DE

Explanation: The NVRAM chip stores user-definable system parameters, also referred to as NVRAM variables or EEPROM parameters. The parameters allow administrators to control

variables such as the default boot device and boot command. The NVRAM also contains writeable areas for user-controlled diagnostics, macros, and device aliases. NVRAM is where the system identification information is stored, such as the host ID, Ethernet address, and time-of-day (TOD) clock.

Examples of NVRAM variables:

Variable Default Description boot-device disk or net The device from which to start up.

diag-device net The diagnostic startup source device.

diag-file Empty string Arguments passed to the startup program in diagnostic mode. diag-switch? false Whether to run in diagnostic mode

NEW QUESTION 208

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

What are the two advantages of following this advice?

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

Answer: AD

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

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

NEW QUESTION 210

A datalink can best be described as .

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

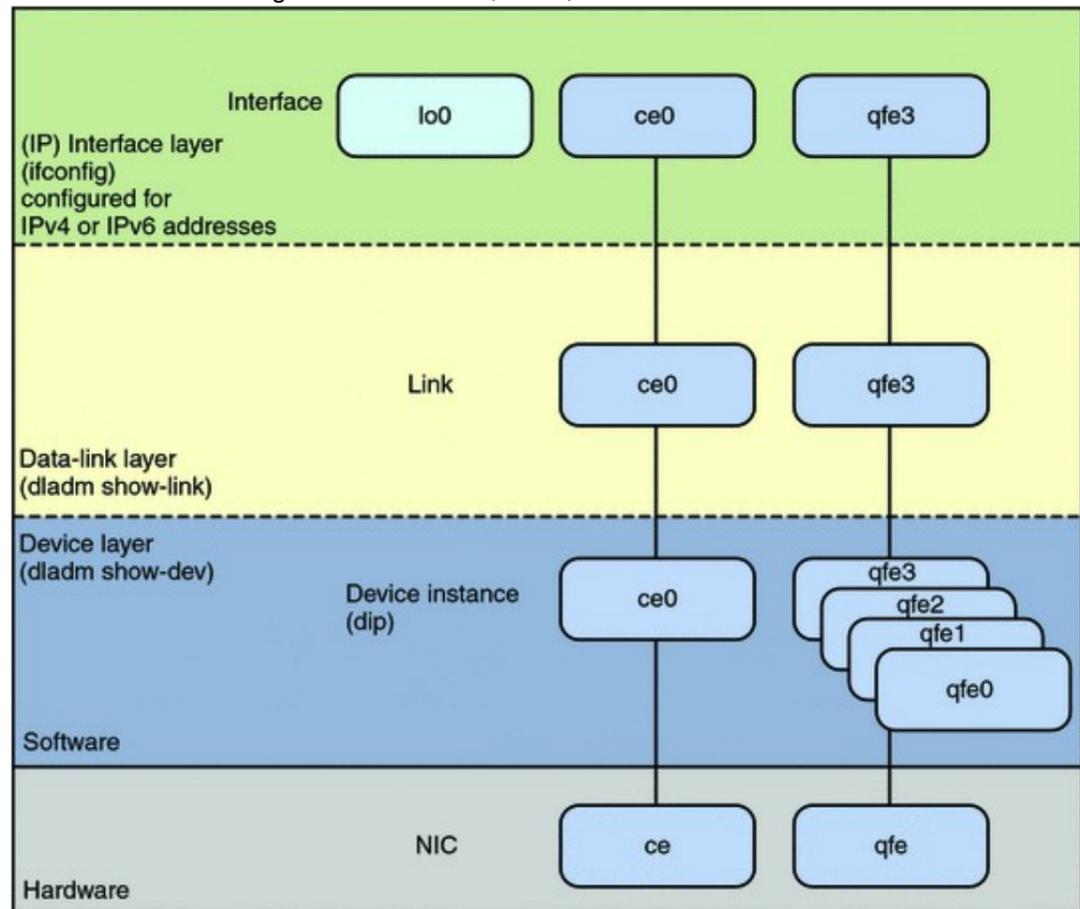
Answer: D

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

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

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

Network Stack Showing Network Devices, Links, and Interfaces:



NEW QUESTION 212

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

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

You then enter: `ipadm create-ip net3`
 What is the output?

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

Answer: B

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

NEW QUESTION 213

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

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

Answer: AD

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

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

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

D: If two client machines need to be installed with the same version of the Oracle Solaris 11 OS but need to be installed differently in other ways, then create two AI manifests for the AI install service. The different AI manifests can specify different packages to install or a different slice as the install target, for example.

NEW QUESTION 218

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

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

Answer: D

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

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

NEW QUESTION 222

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

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

Answer: C

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

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

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

NEW QUESTION 227

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

```
export dennis_says
```

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

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

```
Echo $dennis_says
```

What will he see, and why?

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

Answer: A

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

Note: /etc/skel/local.profile

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

NEW QUESTION 232

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

- A. You can create only one manifest per install service
- B. If you need more than one manifest create multiple install services.
- C. If two client machines have different architectures and need to be installed with the same version of the Oracle Solaris 11 OS, then create two AI manifests and a single install service.
- D. You need a separate install service for each different client architecture that you plan to install, and for each different version of the Oracle Solaris 11 OS that you plan to install on client systems.
- E. If two client machines have different architectures and need to be installed with different versions of the Oracle Solaris 11 OS, then create two AI manifests and two install services.

- F. The install server needs to be able to access an Oracle Solaris Image Packaging System (IPS) software package repository; the clients do not.
 G. The install server can be either an x86 machine or a SPARC machine.

Answer: BEF

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

Note: An AI manifest provides installation instructions.

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

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

NEW QUESTION 233

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 236

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

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

Answer: BEF

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

Mount: mount-point /DS9 does not exist.

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

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

E: showmount

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

showmount [-ade] [hostname]

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

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

F: * mountd Daemon

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

* Commands for Troubleshooting NFS Problems

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

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

NEW QUESTION 239

Which three options accurately describe Oracle Solaris 11 zones?

- A. can be NFS servers
- B. are whole root type only
- C. cannot have their own time zone setting
- D. can execute zfs and zpool commands (from a non-global zone)
- E. are virtualized operating system environments, each with its own dedicated OS and kernel
- F. are virtualized operating system environments, created with a single instance of the OS shared kernel

Answer: ADF

Explanation: A: Zones can use Oracle Solaris 11 products and features such as the following: Oracle Solaris ZFS encryption Network virtualization and QoS CIFS and NFS

C: Non-global zones cannot modify the system clock by default, but each zone can have a separate time zone setting.

F (not E): The Oracle Solaris Zones partitioning technology is used to virtualize operating system services and provide an isolated and secure environment for running applications. A zone is a virtualized operating system environment created within a single instance of the Oracle Solaris operating system.

NEW QUESTION 242

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 247

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 250

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 253

The current ZFS configuration on server is:

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

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

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

You need to backup the /data file system while the file system is active.
 Select the option that creates a full backup of the /data file system and stores the backup on server in the pool named backup.

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

Answer: C

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

NEW QUESTION 255

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

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

Answer: E

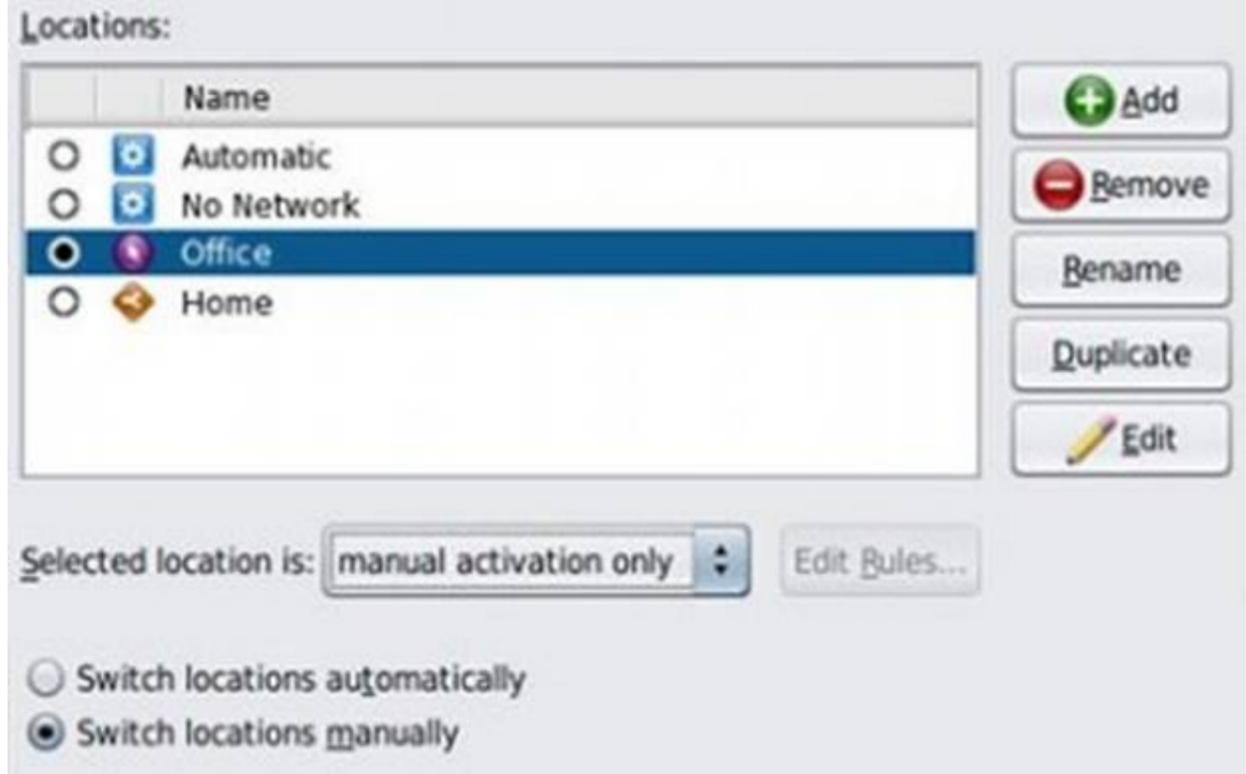
NEW QUESTION 256

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

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

Answer: AD

Explanation: A: Conditional and system locations can be manually activated, which means that the location remains active until explicitly disabled.
 D: A location comprises certain elements of a network configuration, for example a name service and firewall settings, that are applied together, when required. You can create multiple locations for various uses. For example, one location can be used when you are connected at the office by using the company intranet. Another location can be used at home when you are connected to the public Internet by using a wireless access point. Locations can be activated manually or automatically, according to environmental conditions, such as the IP address that is obtained by a network connection.



NEW QUESTION 261

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 262

Which two SMF milestones can be specified at boot time?

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

Answer: AC

Explanation: The milestones that can be specified at boot time are none

single-user multi-user

multi-user-server all

NEW QUESTION 263

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

`pgrep userprogram l5317`

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

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

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

Answer: B

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

To terminate a process:

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

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

NEW QUESTION 268

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

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

Answer: A

NEW QUESTION 270

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 271

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

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

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

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

Answer: G

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

NEW QUESTION 275

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

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

```
Solaris originonlinehttp://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 repositior
- 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.comRefresh 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/releaselIssue 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 276

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

What is true?

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

Answer: C

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

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

Permission 7 full

6 read and write

5 read and execute 4 read only

3 write and execute 2 write only

1 execute only

0 none

0 --- no permission 1 --x execute

2 -w- write

3 -wx write and execute 4 r-- read

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

7 rwx read, write and execut

Solaris: Solaris Advanced User's Guide

NEW QUESTION 281

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

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

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

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

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

Answer: C

Explanation: The following command marks the address object net1/v4a up that was previously marked down.

```
# ipadm up-addr net1/v4a
```

NEW QUESTION 282

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 286

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 289

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 291

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 295

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 298

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 299

.....

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