

1Z0-821 Dumps

Oracle Solaris 11 System Administrator

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



NEW QUESTION 1

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

How can this be accomplished?

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

Answer: E

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

The value assigned by umask is subtracted from the default.

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

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

NEW QUESTION 2

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 file system
- C. Shadow Data Migration
- D. Ability to split a mirrored ZFS storage pool
- E. Ability to use ZFS on the boot drive and boot to a ZFS root file system.
- F. elimination of the swap file system when using ZFS on the root disk

Answer: ABC

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

NEW QUESTION 3

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

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

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

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

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

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

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

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

Answer: C

NEW QUESTION 4

You 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 5

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 flowing output:

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

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

Answer: B

Explanation: Test the remote connection with traceroute.

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

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

NEW QUESTION 6

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 7

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

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

Answer: CD

NEW QUESTION 8

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 9

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

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

Answer: B

Explanation: IPv6 defines three address types: unicast

Identifies an interface of an individual node.

multicast

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

anycast

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

NEW QUESTION 10

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 10

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 13

The advantage of core tiles 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 14

User jack logs in to host Solaris and executes the following command sequence:

```
jack@solaris:~$ cd
jack@solaris:~$ ls -l testfile
-r-xrwxr-- 1 jack other 226 dec 20 20:20 testfile
jack@solaris:~$ id
uid=54326(jack) gid=1(other) groups=1(other)
jack@solaris:~$ id jill
uid=54327(jill) gid=1(other) groups=1(other)
```

Which three statements are correct?

- A. User jack can edit testfile because he has read and write permissions at the group level.
- B. User jack can use cat to output the contents of testfile because he has read permission as the file owner.
- C. User jill can change the permissions of testfile because she has write permission for the file at the group level.
- D. User jill can edit testfile because she has read and write permission at the group level.
- E. User jack can change permissions for testfile because he is the owner of the file.
- F. User jack can change permissions for testfile because he has execute permission for the file.

Answer: DEF

NEW QUESTION 15

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 16

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 20

Review the boot environment information displayed on your system:

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

Which two options accurately describe the newBE boot environment?

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

Answer: BC

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

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

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

You cannot create a snapshot of an unbootable boot environment.

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

You cannot rename an unbootable boot environment. (C)

NEW QUESTION 24

The line

```
set noexec_user_stack=1
```

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

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

Answer: D

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

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

Reboot the system.

```
# reboot
```

NEW QUESTION 26

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 31

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 36

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 38

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 39

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 41

Which two options are characteristics of a fast reboot?

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

Answer: CD

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

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

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

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

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

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

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

NEW QUESTION 44

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

Which option would you choose?

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

Answer: D

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

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

NEW QUESTION 46

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 49

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 54

Given:

`file1` and `file2` are text files. `dir1` and `dir2` are directories.

Which two commands will be successful?

- A. `cp dir1 dir1`
- B. `cp dir1 file1`
- C. `cp file? dir1`
- D. `cp fil`
- E. `dir1`
- F. `cp file% dir2`
- G. `cp file1 file2 dir1`

Answer: CF

Explanation: C: Here the wildcard character `?` is used (Matches any single character). `file1` and `file2` will be copied into `dir1`

F: the two files `file1` and `file2` are copied into directory `dir1`. Note: `cp` - copy files and directories

Copy SOURCE to DEST, or multiple SOURCE(s) to DIRECTORY.

`Cp` has three principal modes of operation. These modes are inferred from the type and count of arguments presented to the program upon invocation.

* When the program has two arguments of path names to files, the program copies the contents of the first file to the second file, creating the second file if necessary.

* When the program has one or more arguments of path names of files and following those an argument of a path to a directory, then the program copies each source file to the destination directory, creating any files not already existing.

* When the program's arguments are the path names to two directories, `cp` copies all files in the source directory to the destination directory, creating any files or directories needed. This mode of operation requires an additional option flag, typically `r`, to indicate the recursive copying of directories. If the destination directory already exists, the source is copied into the destination, while a new directory is created if the destination does not exist.

NEW QUESTION 56

`alice` is a user account used by Alice on a Solaris 11 system. `sadmin` is a role account on the same system.

Your task is to add the command `/usr/sbin/cryptoadm` to the Network management profile, so that Alice can execute it, while assuming the `sadmin` role.

Select the three activities necessary to accomplish this.

- A. To the file `/etc/security/prof_attr`, add the line: `Network Management: solaris:cmd:RO::/usr/sbin/cryptoadm:euid=0`
- B. To the file `/etc/security/auth_attr`, add the line: `Network Management:solaris:cmd:RO::/usr/sbin/cryptoadm:euid=0`
- C. To the file `/etc/security/exec_attr.d/local-entries`, add the line: `Network Management: solaris:cmd:RO::/usr/sbin/cryptoadm:euid=0`
- D. Run the roles `alice` to ensure that `alice` may assume the role `sadmin`.
- E. Run the command `profiles sadmin` to ensure that the role `sadmin` includes the network Management profile.
- F. Run the command `profiles alice` to ensure that the Alice has permissions to access the Network management profile.
- G. Run the command `profiles "Network management"` to ensure that the Network management profile includes the `sadmin` role.

Answer: CDG

Explanation: C: `/etc/security/exec_attr` is a local database that specifies the execution attributes associated with profiles. The `exec_attr` file can be used with other sources for execution profiles, including the `exec_attr` NIS map and NIS+ table.

A profile is a logical grouping of authorizations and commands that is interpreted by a profile shell to form a secure execution environment.

NEW QUESTION 59

Which two statements are true when updating Solaris 11 from one Support Respository Update (SRU) to another SRU by using the `pkg update` command?

- A. By default, the `pkg update` command automatically creates a backup Boot Environment whenever the kernel is affected by the update.

- B. By default, the pkg update command automatically creates a new Boot Environment whenever the kernel is affected by the update.
- C. The pkg update command can only be used to update to a newer SRU.
- D. The pkg update command can be used to update to a newer or older SRU.
- E. By default, the pkg update command always updates Solaris 11 to the first SRU that was released after the Current SRU.
- F. The pkg update command can only be performed while running in the single-user milestone.

Answer: BC

NEW QUESTION 61

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 62

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

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

- A. /etc/security/policy.conf /etc/passwd
- B. /etc/passwd /etc/shadow
- C. /etc/security /etc/passwd
- D. /etc/shadow /etc/passwd

Answer: B

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

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

NEW QUESTION 65

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 68

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 73

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

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

Answer: D

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

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

NEW QUESTION 74

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

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

Answer: B

NEW QUESTION 76

Identify the correct description of an IPS image.

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

Answer: D

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

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

NEW QUESTION 77

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

```

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

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

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

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

```

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

Answer: B

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

For example:

```

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

```

NEW QUESTION 82

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 83

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 84

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 85

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 89

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 91

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 93

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 95

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 99

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 104

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 107

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

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

Answer: ABDG

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

NEW QUESTION 111

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

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

Answer: E

Explanation: The `ipadm delete-addr` subcommand removes addresses from interfaces. To remove an address from the IPMP group, type the following command:
`ipadm delete-addr addrobj`
The `addrobj` uses the naming convention `inder-interface/user-string`.

NEW QUESTION 113

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 116

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 118

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 121

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

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

Answer: C

Explanation: REPOSITORY DESCRIPTION

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

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

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

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

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

NEW QUESTION 122

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

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

Answer: ABC

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

NEW QUESTION 123

A user brian is configured to use the bash shell. His home directory is /export/home/brian, and contains a .profile and a .bashrc file.

In the .profile, there are these lines: genius =ritchie

export genius

In the .bashrc us this line: genius=kernighan

In /etc/profile are these lines: genius=thompson

export genius

When brian logs in and asks for the value of genius, what will he find, and why?

- A. genius will be ritchie, because that was the value exported in .profile.
- B. genius will be kernighan, because .bashrc executes after .profile.
- C. genius will be ritchie because variable settings in .profile take precedence over variable settings in .bashrc.
- D. genius will be ritchie because .profile executes after .bashrc.
- E. genius will be thompson because /etc/profile system settings always override local settings.

Answer: C

NEW QUESTION 127

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 132

Examine this command and its output:

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

```
Name USED AVAIL REFER MOUNTPOINT
```

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

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

Which two conclusions can be drawn based on this output?

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

Answer: AB

NEW QUESTION 135

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

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

Answer: BC

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

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

C: How to Install the solaris10 Branded Zone

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

NEW QUESTION 139

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 140

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 145

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 149

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 151

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

```

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

```

```

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

```

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

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

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

Answer: C

NEW QUESTION 155

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

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

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

Answer: B

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

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

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

In short,

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

NEW QUESTION 159

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 161

Given the following output of the zpool status command:

```
pool: pool1
state: ONLINE
scan: none requested
config:
NAME                STATE      READ    WRITE   CKSUM
pool1               ONLINE    0       0       0
  raidz1-0          ONLINE    0       0       0
    c3t3d0           ONLINE    0       0       0
    c3t4d0           ONLINE    0       0       0
    c3t5d0           ONLINE    0       0       0
    c3t6d0           ONLINE    0       0       0
```

Identify the correct statement regarding pool1's configuration.

- A. Data written to pool1 will be striped across four disk components.
- B. The rsdz1-0 and c3t640 components are submirrors of pool1.
- C. Data will only be striped across the three disks in rsidz configuration.
- D. The configuration is a bug in Solaris 11; it cannot be created by an administrator.

Answer: B

NEW QUESTION 166

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 168

The default publisher on your system is:

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

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

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

Answer: D

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

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

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

Note: Example Adding a Publisher

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

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

Example Specifying the Preferred Publisher

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

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

Example Enabling or Disabling a Publisher

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

You can modify the properties of a disabled publishers.

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

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

NEW QUESTION 173

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

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

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

Answer: A

NEW QUESTION 177

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 180

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 182

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: `kill 15317`

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

- A. You need to run the `kill` 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 `kill` 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 `kill` 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 185

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 client
- B. Configure the repository on the client by using the `svccfg -s` command so that the Solaris publisher is connected to the new repository
- C. Refresh the application/pkg/server service
- D. Issue the `pkgrepo refresh` command to refresh the repository catalog
- E. Configure the publisher on the client using the `svcfg -s` command so that the Solaris publisher is connected to the repository at `http://sysA.example.com` Refresh the application/pkg/server service
- 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.com`
- H. Issue the `pkg update` command to update the OS image.
- I. Add the new publisher `http://sysA.example.com` Solaris Use the `pkg set-publisher` command to set the publisher search order and place `http://sysA.example.com` of `http://pkg.oracle.com/solaris/release` Issue the `pkg publisher` command to view the publisher
- J. Set the new publisher to stick
- K. Issue the `pkg update` command to update the OS image.

Answer: C

Explanation: You can use the `pkg set-publisher` command to change a publisher URI. Changing a Publisher Origin URI

To change the origin URI for a publisher, add the new URI and remove the old URI. Use the `-g` option to add a new origin URI. Use the `-G` option to remove the old origin URI.

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

```
-G http://pkg.example.com/release example.com
```

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

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

NEW QUESTION 187

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

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

Answer: B

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

1. Set up the `/etc/default/login` file with the desired values for `SYSLOG` and `SYSLOG_FAILED_LOGINS`

Edit the `/etc/default/login` file to change the entry. Make sure that `SYSLOG=YES` is uncommented.

2. Create a file with the correct permissions to hold the logging information. Create the authlog file in the /var/adm directory.
3. Edit the syslog.conf file to log failed password attempts. Send the failures to the authlog file.

NEW QUESTION 191

Which command would you use from the bash shell to determine the total amount of physical memory installed in your Solaris system (x86 and SPARC)?

- A. uname -a
- B. prtconf | grep -i memory
- C. sysdef | grep -i memory
- D. vmstat
- E. prtdiag | grep -i memory

Answer: B

Explanation: The prtconf command prints the system configuration information. The output includes the total amount of memory, and the configuration of system peripherals formatted as a device tree.

If a device path is specified on the command line for those command options that can take a device path, prtconf will only display information for that device node.

NEW QUESTION 192

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

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

Answer: C

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

Example:

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

NEW QUESTION 196

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 197

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 199

You are going to create live zones on your server. Disk space is critical on this server so you need to reduce the amount of disk space required for these zones. Much of the data required for each of these zones is identical, so you want to eliminate the duplicate copies of data and store only data that is unique to each zone. Which two options provide a solution for eliminating the duplicate copies of data that is common between all of these zones?

- A. Create the zones by using sparse root zones.
- B. Set the dedup property to on and the dedupratio to at least 1.5 for the zpool. Create a separate ZFS file system for each zone in the zpool.
- C. Put all of the zones in the same ZFS file system and set the dedupratio property for the ZFS file system to at least 1.5.
- D. Put all of the zones in the same ZFS file system and set the dedup property for the file system to on.
- E. Put each zone in a separate ZFS file system within the same zpool.
- F. Set the dedup property to on for each ZFS file system.

Answer: DE

Explanation: In Oracle Solaris 11, you can use the deduplication (dedup) property to remove redundant data from your ZFS file systems. If a file system has the dedup property enabled, duplicate data blocks are removed synchronously. The result is that only unique data is stored, and common components are shared between files.

NEW QUESTION 200

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

- A. ppgsz
- B. pargs
- C. pmap
- D. pgrep

Answer: B

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

NEW QUESTION 204

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

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

201301(Pr201301)

What are the two problems with your plan?

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

Answer: CD

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

D: Group names contain only lowercase characters and numbers.

NEW QUESTION 207

You have installed software updates to a new boot environment (BE) and have activated that the booting to the new BE, you notice system errors. You want to boot to the last known good configuration.

Which option would you use on a SPARC system to boot to the currentBE boot environment?

- A. boot -L currentBE
- B. boot -Z rpool/ROOT/currentBE
- C. boot -a Enter the currentBE dataset name when prompted.
- D. boot rpool/ROOT/currentBE
- E. boot -m currentBE
- F. beadm activate currentBE

Answer: F

Explanation: You can change an inactive boot environment into an active boot environment. Only one boot environment can be active at a time. The newly activated boot environment becomes the default environment upon reboot.

How to Activate an Existing Boot Environment

1. Use the following command to activate an existing, inactive boot environment: beadm activate beName

beName is a variable for the name of the boot environment to be activated. Note the following specifications.

beadm activate beName activates a boot environment by setting the bootable pool property, bootfs, to the value of the ROOT dataset of the boot environment that is being activated.

beadm activate sets the newly activated boot environment as the default in the menu.lst file.

2. Reboot.

The newly activated boot environment is now the default on the x86 GRUB menu or SPARC boot menu.

NEW QUESTION 210

You are troubleshooting network throughput on your server.

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

The correct command is .

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

Answer: B

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

NEW QUESTION 214

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

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

- A. pkgchk opendap
- B. pkginfo opendap
- C. pkg contents opendap
- D. pkg verify opendap
- E. pkg set-property signature-policy verify

Answer: A

Explanation: pkgchk checks the accuracy of installed files or, by using the -l option, displays information about package files. pkgchk checks the integrity of directory structures and files. Discrepancies are written to standard error along with a detailed explanation of the problem.

NEW QUESTION 217

Which three options describe the purpose of the zonep2vchk command?

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

Answer: CDE

Explanation: zonep2vchk

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

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

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

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

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

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

for p2v into a Solaris 10 Branded zone.

NEW QUESTION 220

Review the boot environments displayed on your system:

BE	Active	Mountpoint	Space	Policy	Created
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 option describes the solaris-1 BE?

- A. It is active on the next reboot.
- B. It is active now.
- C. It is inactive.
- D. It is unbootable.
- E. It is active now and on reboot.
- F. It has been removed and will no longer be available after the next reboot.

Answer: E

Explanation: In the below output, NR (now running) means the BE is active now and will be the active BE on reboot.

Example:

Display your existing BE information.

beadm list

BE Active Mountpoint Space Policy Created

```
-----
solaris NR / 12.24G static 2011-10-04 09:42
```

NEW QUESTION 222

View the Exhibit and review the disk configuration.

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

Primary label contents:

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

format> █
```

The following command is executed on the disk: zpool create pool1 c3t3d0s0

What is the result of executing this command?

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

Answer: B

NEW QUESTION 227

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

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

Answer: ABDFH

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

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

```
$ sudo pkg install slim_install
```

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

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

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

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

NEW QUESTION 229

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

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

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

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

Response:

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

What is the next logical troubleshooting action?

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

Answer: D

Explanation: Check the physical connection.

NEW QUESTION 234

You are currently working in both your home directory and the system directory `/tmp`. You are switch back and forth with full path names. Which pair of `cd` commands will provide you with a shortcut to switch between these two locations?

- A. `cd ~` and `cd -`
- B. `cd` and `cd`.
- C. `cd ~` and `cd`
- D. `cd *` and `cd .`

Answer: A

Explanation: In the Bourne Again, C, Korn, TC, and Z shells, the tilde (`~`) is used as a shortcut for specifying your home directory.

```
cd -
```

It's the command-line equivalent of the back button (takes you to the previous directory you were in).

Note:

To make certain that you are in your home directory, type the `cd` (change directory) command. This command moves you to your home (default) directory.

NEW QUESTION 237

You are creating a non-global zone on your system.

Which option assigns a zpool to a non-global zone, and gives the zone administrator permission to create zfs file system in that zpool?

- A. While creating the non-global zone, make the following entry: `add deviceset match=/dev/rdisk/c4t0d0endBoot` the zone and log in the zone as root
- B. Create the zpool: `zpool create pool2 c4t0d0`In the non-global zone, root can now create ZFS file system in the pool2 zpool
- C. In the global zone, create the zpool: `global# zpool create pool2 c4t1d0`While creating the no-global zone, make the following entry: `add datasetset name=pool2endadd fsset dir=pool1set special=pool1set type=zfspool1endBoot` the zone, log in the zone as root, and create the zfs file system in the pool2 zpool.
- D. In the global zone, create the zpool:`global#zpool create pool2 c4t1d0`While creating the global zone, make the following entry: `add datasetset name=pool2endBoot` the zone, log in to the zone as root and create the zfs file systems in the pool2 zpool.
- E. In the global zone, create the zpool and the ZFS file systems that you want to use in the non-global zone: `global#zpool create pool2 c4t1d0global#zfs create pool2/data`While creating the non-global zone, make the following entry for each ZFS file system that you want to make available in the zone: `add fsset dir=/dataset special=pool2/dataset type=zfsend`
- F. Create the zpool in the global zone: `global#zpool create pool2 c4t1d0`Boot the non- global zone, log in to the zone as root, and issue this command to delegate ZFS permissions to root: `non-global zone# zfs allow root create , destroy, mount pool2`Log in to the non-global zone create ZFS file systems in the pool2 zpool.

Answer: C

Explanation: <http://docs.oracle.com/cd/E19253-01/819-5461/gbbst/index.html>

NEW QUESTION 240

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

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

Identify two correct statements.

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

Answer: AB

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

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

```
$ su - root
```

```
Password: Type root password
```

```
#
```

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

NEW QUESTION 242

United States of America export laws include restrictions on cryptography.

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

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

Answer: BC

Explanation: B: Binary Signatures for Third-Party Software

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

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

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

NEW QUESTION 246

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

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

Which three statements are correct?

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

Answer: ACD

Explanation: A: dir followed by a single letter.

C: dir followed by any characters ending with a.

D: dir followed by any characters, then character b, then one single character. only dirabc matches

NEW QUESTION 247

The OpenBoot firmware controls the operation of the system before the operating system is loaded.

Which four tasks are directly controlled by the OpenBoot firmware?

- A. Provides a list of boot entries from which to choose
- B. Allows hardware to identify itself and load its own plug-in device driver
- C. Loads the boot loader from the configured boot device
- D. Performs basic hardware testing

- E. Installs the console
- F. Reads and executes the boot archive
- G. Extract and executes the kernel image

Answer: ABCD

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

Some notable features of OpenBoot firmware.

- * Plug-in Device Drivers (B)

Most common tasks that you perform using OpenBoot

- * Booting Your System

The most important function of OpenBoot firmware is to boot the system.

Note: If auto-boot? is true, the system will boot from either the default boot device or from the diagnostic boot device depending on whether OpenBoot is in diagnostic mode. (C)

NEW QUESTION 249

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

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

Which four correctly describe the output?

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

Answer: ACEG

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

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

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

Note: Identifying ZFS Snapshot Differences (zfs diff)

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

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

File or Directory Change Identifier

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

NEW QUESTION 250

You want to display network interface information. Which command should you use?

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

Answer: A

NEW QUESTION 251

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

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

Answer: AF

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

NEW QUESTION 256

After installing the OS, you boot the system and notice that the syslogd daemon is not accepting messages from remote systems. Which two options should you select to modify the syslogd daemon configuration so that it accepts messages from remote systems?

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

Answer: BD

Explanation: B: The `/etc/default/syslogd` file contains the following default parameter settings. See FILES.
`LOG_FROM_REMOTE`
Specifies whether remote messages are logged. `LOG_FROM_REMOTE=NO` is equivalent to the `-t` command-line option. The default value for `LOG_FROM_REMOTE` is YES.

NEW QUESTION 260

Identify the two security features incorporated in the Oracle Solaris 11 Cryptographic Framework.

- A. Layer 5 IP address encryptions
- B. Internet protocol security
- C. Diffie-Kerberos coaxial key encryption
- D. Signed cryptographic plugins (providers)
- E. Kernel support for signed antivirus plugins

Answer: DE

Explanation: The framework enables providers of cryptographic services to have their services used by many consumers in the Oracle Solaris operating system. Another name for providers is plugins. The framework allows three types of plugins:

* User-level plugins - Shared objects that provide services by using PKCS #11 libraries, such as `pkcs11_softtoken.so.1`.

* Kernel-level plugins - Kernel modules that provide implementations of cryptographic algorithms in software, such as AES.

Many of the algorithms in the framework are optimized for x86 with the SSE2 instruction set and for SPARC hardware.

* Hardware plugins - Device drivers and their associated hardware accelerators. The Niagara chips, the `ncp` and `n2cp` device drivers, are one example. A hardware accelerator offloads expensive cryptographic functions from the operating system. The Sun Crypto Accelerator 6000 board is one example.

NEW QUESTION 261

Examine this command and its output:

```
# zoneadm list Global
zone2
```

Which two outcomes can be deduced from this output?

- A. There is exactly one nonglobal zone installed.
- B. There is one nonglobal zone running.
- C. There is at least one oneglobal zone configured.
- D. There is one nonglobal zone that is not configured.
- E. There is one nonglobal zone that is not running.
- F. There is one nonglobal zone that is not installed.
- G. There is exactly one nonglobal zone configured.

Answer: BC

NEW QUESTION 264

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

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

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

Answer: AEF

NEW QUESTION 269

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

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

These are the steps that you have followed:

1. Download and rsync the contents of the Oracle Solaris11 repository ISO image to the /repo directory.
2. Configure the repository server service properties. The svcprop command display, the IPS related properties:
pkg/inst_root astring/repo pkg/readonly Boolean true
The ls command displays the contents of the /repo directory:
#ls/repo
Pkg5.repository publisher
The svcs publisher command shows the svc: /application/pkg/server: default service is online.
The pkg publisher command shows the svc: /application/pkg/server: default service is online.
The pkg publisher command still displays: PUBLISHERTYPESTATUSURI
solarisoriginonlinehttp://pkg.oracle.com/solaris/release/
Which steps needs to be performed to set the local IPS publisher to the local IPS repository/repo?

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

Answer: C

Explanation: Set the Publisher Origin To the File Repository URI

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

Example:

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

NEW QUESTION 274

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

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

Answer: D

Explanation: Searching for Packages

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

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

pkg search

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

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

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

NEW QUESTION 275

Your server has one zone named dbzone (hat has been configured, but not yet installed). Which command would you use to view all the options that were used to configure this zone?

- A. zoneadm list -icv dbzone
- B. zones tat -c summary dbzone
- C. zonecfg -z dbzone info
- D. zonecfg -icv dbzone info

Answer: C

Explanation: zonecfg info

Display information about the current configuration. If resource-type is specified, displays only information about resources of the relevant type. If any property-name value pairs are specified, displays only information about resources meeting the given criteria. In the resource scope, any arguments are ignored, and info displays information about the resource which is currently being added or modified.

Note: zonecfg -z

zonename. Specify the name of a zone. Zone names are case sensitive. Zone names must begin with an alphanumeric character and can contain alphanumeric characters, the underscore (_) the hyphen (-), and the dot (.). The name global and all names beginning with SUNW are reserved and cannot be used.

Incorrect Answer

A: The zoneadm utility is used to administer system zones. A zone is an application container that is maintained by the operating system runtime.

list option:

Display the name of the current zones, or the specified zone if indicated. B: No such command.

D: no such options zonecfg -icv

NEW QUESTION 280

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

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

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

Answer: BE

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

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

NEW QUESTION 285

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

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

Answer: B

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

pkg info example:

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

Summary: The p7zip compression and archiving utility

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

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

Build Release: 5.11

Branch: 0.175.0.0.0.2.537

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

Size: 6.73 MB

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

NEW QUESTION 286

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

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

Answer: ABC

Explanation: How to Monitor All Failed Login Attempts

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

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

NEW QUESTION 289

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

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

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

Answer: A

NEW QUESTION 292

Subnets are created by using .

- A. subnet
- B. netmask
- C. unicast
- D. broadcast

Answer: B

Explanation: The process of subnetting involves the separation of the network and subnet portion of an address from the host identifier. This is performed by a bitwise AND operation between the IP address and the (sub)network prefix. The result yields the network address or prefix, and the remainder is the host identifier. 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 295

Which three statements are true concerning Image Packaging System (IPS) incorporation package?

- A. Installing an incorporation package does not install any other packages.
- B. Every feature or tool has a separate IPS incorporation.
- C. They constrain the versions of packages they incorporate.
- D. They are a content management tool and not a version management tool.
- E. Their dependencies are always of TYPE-REQUIRE.
- F. They are defined by their manifest

Answer: ACE

NEW QUESTION 298

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 1Z0-821 Exam with Our Prep Materials Via below:

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