



Oracle

Exam Questions 1Z0-821

Oracle Solaris 11 System Administrator

NEW QUESTION 1

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.

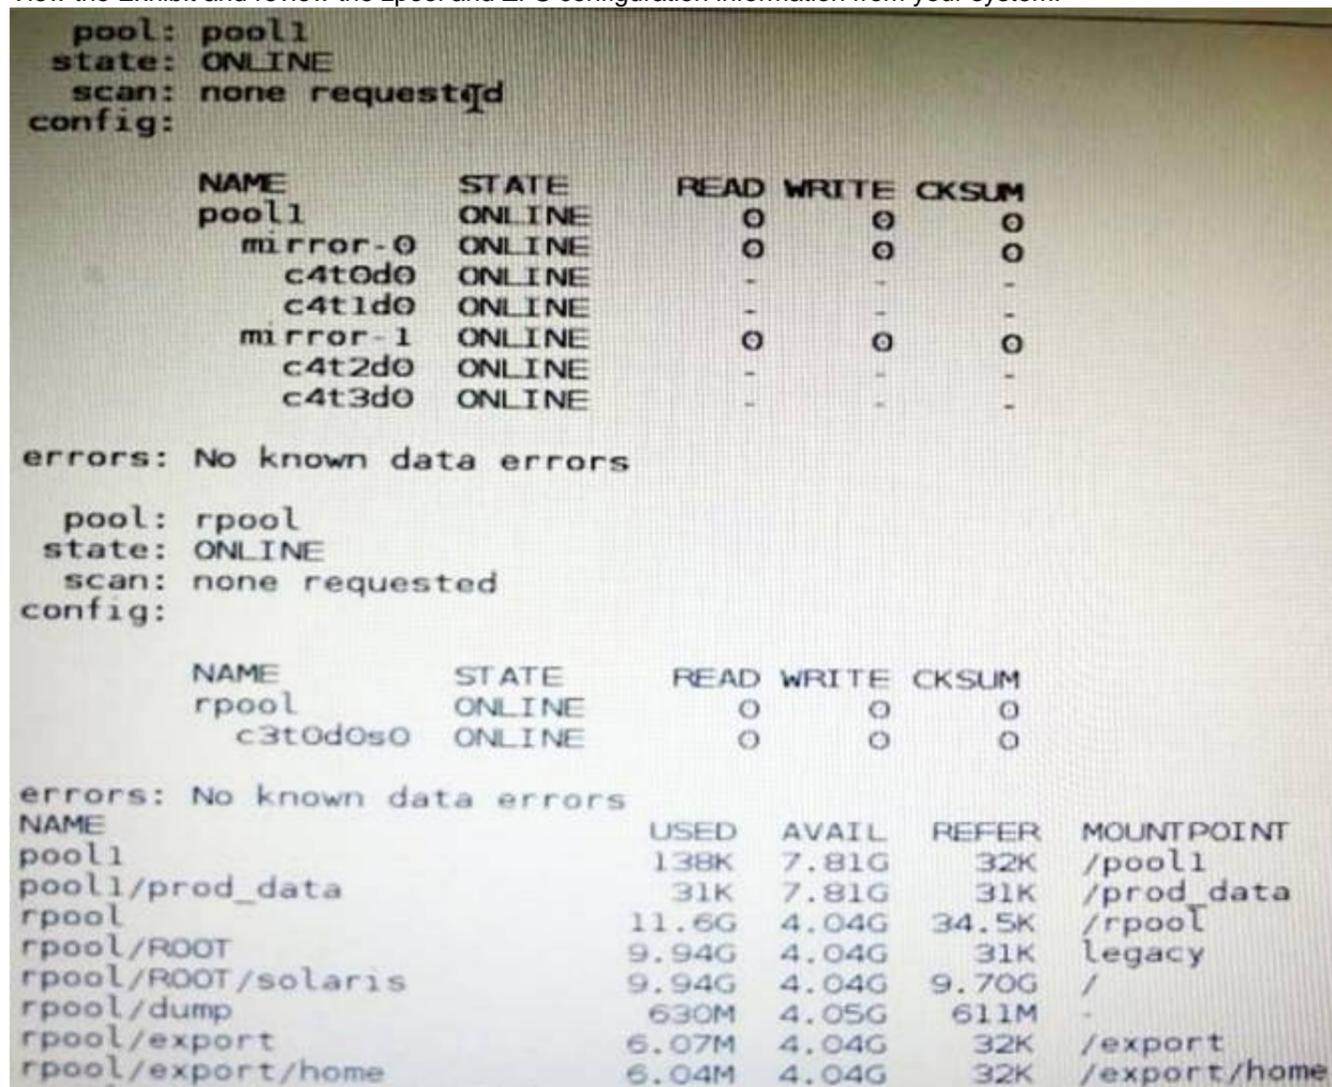
beadm list

BE Active Mountpoint Space Policy Created

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

NEW QUESTION 2

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



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

- A. zpool split pool1 pool2 c4t0d0 c4t2d0 zpool import pool2 zfs set mountpoint = /dev_data pool2/prod_data
- B. zpool detach pool1 pool2 zpool attach pool2 zfs set mountpoint = /dev_data pool2/prod_data
- C. zpool split pool1/prod_data -n pool2/dev_data zfs set mountpoint = /dev_data pool2/prod_data
- D. zpool split pool1 pool2 c4t0d0 c4t2d0 zpool import pool2

Answer: D

Explanation:

In this Solaris release, you can use the zpool split command to split a mirrored storage pool, which detaches a disk or disks in the original mirrored pool to create another identical pool. After the split operation, import the new pool.

NEW QUESTION 3

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 4

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 5

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 6

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

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

D. two etherstubs

Answer: A

NEW QUESTION 7

The current ZFS configuration on server is:

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

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

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 8

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 9

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 10

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 10

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

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

Answer: BEF

Explanation:

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

Mount: mount-point /DS9 does not exist.

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

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

E: showmount

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

showmount [-ade] [hostname]

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

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

F: * mountd Daemon

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

* Commands for Troubleshooting NFS Problems

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

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

NEW QUESTION 15

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 17

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 19

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

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

Answer: BEF

Explanation:

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

slice as the install target, for example.

Note: An AI manifest provides installation instructions.

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

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

NEW QUESTION 20

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

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

Answer: C

Explanation:

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

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

```
# ipadm show-addr
```

```
ADDROBJ TYPE STATE ADDR
```

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

NEW QUESTION 21

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

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

Answer: D

Explanation:

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

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

NEW QUESTION 25

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 30

user1, while in his home directory, is attempting to run the following command in his home directory: cp bigfile verybig
The system displays the following error:
cp: cannot create verybig: Disc quota exceeded
Your initial troubleshooting shows that the df -h command indicates he is at 100% capacity. What command would you use to increase the disk space available to the user?

- A. zfs get quota rpool/export/home/user1
- B. zfs userused@user1
- C. zfs quota=none /rpool/export/home/user1
- D. df -h | grep user1
- E. zfs set quota=none /rpool/export/home/user1

Answer: E

Explanation:

ZFS quotas can be set and displayed by using the zfs set and zfs get commands. We can remove the quota restriction by setting to quota to none.

NEW QUESTION 34

Which two are true about accounts, groups, and roles in the Solaris user database?

- A. All Solaris user accounts must have a unique UID number.
- B. A Solaris account name may be any alphanumeric string, and can have a maximum length of 8 characters.
- C. Account UID numbers 0-09 are system-reserved.
- D. The GID for an account determines the default group ownership of new files created by that account.
- E. The groups that an account is a member of are determined by the entries in the/etc/group file.

Answer: AB

Explanation:

A: Solaris uses a UID (User ID) to identify each user account. The UID is a unique number assigned to each user. It is usually assigned by the operating system when the account is created.

B: In Solaris the account name can include any alphanumeric string (and . _ -). The maximum length is 8 characters.

NEW QUESTION 37

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 42

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 44

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 45

Identify the Automated Installer's (AI) equivalent to jumpStart's finish scripts and sysidcfg files.

- A. Manifest files
- B. SMF system configuration profile files
- C. Installadm create - client
- D. IPS software package repository
- E. installadm create-service
- F. svccfg - s application/pkg/server setprop sysidcfg

Answer: B

Explanation:

Comparing sysidcfg File Keywords to System Configuration Profile Directives

The following table compares sysidcfg file keywords with example AI system configuration profile specifications.

sysidcfg File Keyword

System Configuration Profile Directives Etc.

NEW QUESTION 49

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 54

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

- A. printenv
- B. banner

- C. .version
- D. set-env
- E. show-devs
- F. value version

Answer: BC

Explanation:

B: banner

Displays power-on banner.

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

SPARCstation 2, Type 4 Keyboard

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

Displays version and date of the boot PROM.

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

NEW QUESTION 58

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 63

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 66

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

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

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

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

Answer: AC

Explanation:

A: %PART

The amount of cpu used as a percentage of the total cpu in a processor-set to which the zone is bound. A zone can only have processes bound to multiple processor sets if it is the global zone, or if psrset(1m) psets are used. If multiple binding are found for a zone, it's

%PART is the fraction used of all bound psets. For [total] and [system], %PART is the percent used of all cpus on the system.

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

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

NEW QUESTION 71

You are using AI to install a new system. You have added to following information to the AI manifest:
`<configuration type= "zone" name= "dbzone"`
`source = "http://sysA.example.com/zone_cfg/zone.cfg"/>` Which statement is true with regard to the zone.cfg?

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

Answer: A

Explanation:

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

NEW QUESTION 73

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

- A)

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

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

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

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

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

Answer: B

Explanation:

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

For example:

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

NEW QUESTION 75

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 79

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

- A. TCP
- B. UDP
- C. IP

- D. ICMP
- E. Ethernet

Answer: C

Explanation:

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

NEW QUESTION 84

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 88

The core dump configuration for your system is:

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

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

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

The server host name is: zeus

What will the per-process core file be named?

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

Answer: C

Explanation:

Note the first line:

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

The program name is bash The runtime process ID is 2663

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

NEW QUESTION 90

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

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

You want to use the flash archive, located On" /net/bactup_servers/10-system.flar, to install the Operating system in the s10-zone zone.

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

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

Answer: A

Explanation:

The zoneadm command is the primary tool used to install and administer non-global zones. Operations using the zoneadm command must be run from the global

zone on the target system.

How to Install the solaris10 Branded Zone

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

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

NEW QUESTION 92

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 93

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 96

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 99

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 101

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

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

Answer: AC

Explanation:

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

NEW QUESTION 106

You have already generated a 256-bit AES raw key and named the keystore file /mykey. You need to use the key to create an encrypted file system. Which command should you use to create a ZFS encrypted file system named pool1/encrypt using the /mykey keystore?

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

Answer: B

Explanation:

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

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

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

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

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

NEW QUESTION 108

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

- A. It converts an Oracle Solaris 11 host into a SCSI target device that can be accessed over a storage network by Linux, Mac OS, or Windows client systems.
- B. iSCSI targets cannot be configured as dump devices.
- C. It provides support for iSCSI devices that use SLP.
- D. It is used to connect to Fibre Channel or iSCSI Storage Area Network (SAN) environments.
- E. It provides an upgrade and update path to convert your iSCSI LUNs from Solaris 10 systems.

Answer: AB

Explanation:

A: You can configure Common Multiprotocol SCSI TARget, or COMSTAR, a software

framework that enables you to convert any Oracle Solaris 11 host into a SCSI target device that can be accessed over a storage network by initiator hosts.

This means you can make storage devices on a system available to Linux, Mac OS, or Windows client systems as if they were local storage devices. Supported storage protocols are iSCSI, FC, iSER, and SRP.

B: iSCSI targets cannot be configured as dump devices.

NEW QUESTION 110

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

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 114

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 118

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

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

Answer: A

Explanation:

Network configuration in Oracle Solaris 11 includes

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

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

NEW QUESTION 122

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 123

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

To accomplish this, you will create .

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

Answer: D

Explanation:

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

NEW QUESTION 125

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 130

View the following information for a software package:

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

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

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

Answer: B

Explanation:

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

Example:

Listing Information About an Uninstalled Package

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

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

```
Version: 5.2.17
```

```
Build Release: 5.11
```

```
Branch: 0.175.0.0.0.1.530
```

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

```
Size: 44.47 MB
```

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

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

installed packages in the image.

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

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

NEW QUESTION 134

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 138

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

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

Answer: A

Explanation:

Part of the job of cleaning up heavily loaded file systems involves locating and removing files that have not been used recently. You can locate unused files by using the ls or find commands.

Other ways to conserve disk space include emptying temporary directories such as the directories located in /var/tmp or /var/spool, and deleting core and crash dump files.

Note: Core files are generated when a process or application terminates abnormally. Core files are managed with the coreadm command.

For example, you can use the coreadm command to configure a system so that all process core files are placed in a single system directory. This means it is easier to track problems by examining the core files in a specific directory whenever a process or daemon terminates abnormally.

NEW QUESTION 140

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

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

You then ran: ipadm show-if

The result indicated that the interface was down.

You then ran:

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

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

The last command indicated that the interface was up.

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

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

Answer: B

Explanation:

The first IP address is already in use.

NEW QUESTION 141

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 144

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

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

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

Answer: B

Explanation:

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

TIME_EXCEEDED response from each gateway along the path to some host.

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

NEW QUESTION 148

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 149

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 153

View the Exhibit.

```

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

```

Which is true regarding the disk drive?

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

Answer: A

Explanation:

Installing a ZFS Root Pool

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

NEW QUESTION 155

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

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

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

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

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

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

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

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

Answer: C

NEW QUESTION 156

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

NAME (PUBLISHER)VERSIONINFO

Compress/zip3.1.2-0.175.0.0.0.0.537if-

NAMEVERSIONDATECOMMENT

Compress/zip3.109 Dec 2011 04:50:38 ESTNone

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 161

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

- A. You can perform a text or LiveCD installation locally or over the network.

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

Answer: ABDFH

Explanation:

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

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

```
$ sudo pkg install slim_install
```

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

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

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

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

NEW QUESTION 163

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

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

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

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

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

Answer: A

Explanation:

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

Note: zpool attach [-f] pool device new_device

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

NEW QUESTION 166

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 tranceroute -n server.
- F. On Server B, run tranceroute -n servera and tranceroute -n serverc.

Answer: D

Explanation:

Check the physical connection.

NEW QUESTION 169

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 roo
- B. Create the zpool: zpool create pool2 c4t0d0In 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 c4t1d0While 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 c4t1d0While 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/dataWhile 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 c4t1d0Boot 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 pool2Log 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 171

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 172

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 173

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 178

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 182

You want the system to generate an email notification each time one of the services has changed its state. Which option would send an email message to the system administrator whenever a service changes to the maintenance state?

- A. Use the setsc command in ALOM to enable the mail alerts to be sent to a specified email address whenever the fault management facility detects a service change to the maintenance state.
- B. Make an entry in the /etc/syslog.conf file to instruct syslogd to send an email alert when it receives a message from the SMF facility that a service has changed to the maintenance state.
- C. Use the svccfg setnotify command to create a notification and send an email when a service enters the maintenance state.
- D. Use the scvadm command to enable the notification service
- E. Set the -g maintenance option on the netnotify service to send an email when a service enters the maintenance state.

Answer: C

Explanation:

This procedure causes the system to generate an email notification each time one of the services or a selected service has a change in state. You can choose to use either SMTP or SNMP. Normally, you would only select SNMP if you already have SNMP configured for some other reason.

By default, SNMP traps are sent on maintenance transitions. If you use SNMP for monitoring, you can configure additional traps for other state transitions.

1. Become an administrator or assume a role that includes the Service Management rights profile.
2. Set notification parameters.

Example:

The following command creates a notification that sends email when transactions go into the maintenance state

```
# /usr/sbin/svccfg setnotify -g maintenance mailto:sysadmins@example.com
```

NEW QUESTION 185

Review the storage pool information:

```
pool: pool1
state: DEGRADED
status: One or more devices could not be opened. Sufficient replicas exist for
the pool to continue functioning in a degraded state.
action: Attach the missing device and online it using 'zpool online'.
see: http://www.sun.com/msg/ZFS-8000-2Q
scan: none requested
config:
NAME          STATE      READ    WRITE   CKSUM
pool1         DEGRADED   0       0       0
  mirror-0    DEGRADED   0       0       0
    c3t3d0    UNAVAIL   0       0       0 cannot open
    c3t4d0    ONLINE   0       0       0
```

Choose the correct procedure to repair this storage pool.

- A. Shut the system down, replace disk c3t3d0, and boot the system
- B. When the system is booted, execute the zpool clear pool1 command.
- C. Shut the system down, replace disk c3t3d0, and boot the system
- D. When the system is booted execute the zpool online pool1 command.
- E. Shut the system down, replace disk c3t3d0, and boot the system
- F. When the system is booted, execute the zpool replace pool1 c3t3d0 command.
- G. Shut the system down, replace disk c3t3d0, and boot the system
- H. When the system is booted, execute the zpool replace pool1 c3t3d0 c3t3d0 command.

Answer: C

Explanation:

You might need to replace a disk in the root pool for the following reasons: The root pool is too small and you want to replace it with a larger disk

The root pool disk is failing. In a non-redundant pool, if the disk is failing so that the system won't boot, you'll need to boot from an alternate media, such as a CD or the network, before you replace the root pool disk.

In a mirrored root pool configuration, you might be able to attempt a disk replacement without having to boot from alternate media. You can replace a failed disk by using the zpool replace command.

Some hardware requires that you offline and unconfigure a disk before attempting the zpool replace operation to replace a failed disk.

For example:

```
# zpool offline rpool c1t0d0s0
# cfgadm -c unconfigure c1::dsk/c1t0d0
<Physically remove failed disk c1t0d0>
<Physically insert replacement disk c1t0d0>
# cfgadm -c configure c1::dsk/c1t0d0
# zpool replace rpool c1t0d0s0
# zpool online rpool c1t0d0s0
# zpool status rpool
<Let disk resilver before installing the boot blocks>
SPARC# installboot -F zfs /usr/platform/`uname -i`/lib/fs/zfs/bootblk /dev/rdisk/c1t0d0s0 x86# installgrub /boot/grub/stage1 /boot/grub/stage2 /dev/rdisk/c1t9d0s0
```

NEW QUESTION 190

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 194

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 196

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

ADDROBJ	TYPE	STATE	ADDR
lo0/v4	static	ok	127.0.0.1/8
lo0/v6	static	ok	::1/128

Which option describes the state of this server?

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

Answer: C

Explanation:

Only the loopback addresses are configured. No IP address is configured.

NEW QUESTION 200

View the exhibit.

ADDROBJ	TYPE	STATE	ADDR
lo0/v4	static	ok	127.0.0.1/8
net0/_b	dhcp	ok	10.0.2.15/24
net1/_b	dhcp	ok	10.0.3.15/24
lo0/v6	static	ok	::1/128
net0/_a	addrconf	ok	fe80::a00:27ff:fee5:38b9/10
net1/_a	addrconf	ok	fe80::a00:27ff:fe2b:498a/10

The configuration information in the exhibit is displayed on your system immediately after installing the OS. Choose the option that describes the selection made during the installation of the OS to obtain this configuration.

- A. The automatic network configuration option was chosen during the installation of the OS.
- B. The manual network configuration option was chosen during the installation of the OS.
- C. The network was not configured during the installation of the OS.
- D. The DHCP network configuration option was chosen during the installation of the OS.

Answer: A

Explanation:

There are two ways to configure the network configuration: automatic or manual.

In the exhibit we see that DHCP has been used. This indicates an automatic network configuration.

NEW QUESTION 205

You want to display the IP address assignments of the network interfaces. Which command should you use?

- A. ipadm show-if
- B. ipadm show-addr

- C. ipadm show-prop
- D. ipadm show-addrprop

Answer: B

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 207

Review the storage pool information:

```
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

Which statement describes the status of this storage pool?

- A. It is a RAIDZ storage pool and can withstand a single disk failure; data will be striped at: disk components.
- B. It is a double-parity RAIDZ storage pool and can withstand two disk failures; data will be striped across four disk components.
- C. It is an improperly configured RAIDZ storage pool; data will be striped across four disk components, but only three drives are protected with redundancy.
- D. It is an improperly configured RAIDZ storage pool; data will be striped across three disk components, but only three drives are protected with redundancy.

Answer: D

Explanation:

Device c3t6d0 is not included in the RAIDZ storage pool. The other three devices are included in the raidz pool. The data on these devices are protected.
 Note: In addition to a mirrored storage pool configuration, ZFS provides a RAID-Z configuration with either single, double, or triple parity fault tolerance. Single-parity RAID-Z (raidz or raidz1) is similar to RAID-5. Double-parity RAID-Z (raidz2) is similar to RAID-6.

NEW QUESTION 211

Select two correct statements about the authentication services available in Oracle Solaris 11.

- A. Pluggable Authentication Modules (PAM) is used to control the operation of services such console logins and ftp.
- B. The Secure Shell can be configured to allow logins across a network to remote servers without transmitting passwords across the network.
- C. Secure Remote Procedure Calls (Secure RPC) provides a mechanism to encrypt data on any IP Socket connection.
- D. Pluggable Authentication Modules (PAM) is used to implement the Secure Shell in Oracle Solaris 11.
- E. Simple Authentication and Security Layer (SASL) provides a mechanism to authenticate and encrypt access to local file system data.

Answer: AE

Explanation:

A: Pluggable Authentication Modules (PAM) are an integral part of the authentication mechanism for the Solaris. PAM provides system administrators with the ability and flexibility to choose any authentication service available on a system to perform end-user authentication. By using PAM, applications can perform authentication regardless of what authentication method is defined by the system administrator for the given client. PAM enables system administrators to deploy the appropriate authentication mechanism for each service throughout the network. System administrators can also select one or multiple authentication technologies without modifying applications or utilities. PAM insulates application developers from evolutionary improvements to authentication technologies, while at the same time allowing deployed applications to use those improvements. PAM employs run-time pluggable modules to provide authentication for system entry services.
 E: The Simple Authentication and Security Layer (SASL) is a method for adding authentication support to connection-based protocols. Simple Authentication and Security Layer (SASL) is a framework for authentication and data security in Internet protocols. It decouples authentication mechanisms from application protocols, in theory allowing any authentication mechanism supported by SASL to be used in any application protocol that uses SASL. Authentication mechanisms can also support proxy authorization, a facility allowing one user to assume the identity of another. They can also provide a data security layer offering data integrity and data confidentiality services. DIGEST-MD5 provides an example of mechanisms which can provide a data-security layer. Application protocols that support SASL typically also support Transport Layer Security (TLS) to complement the services offered by SASL.

NEW QUESTION 215

Select two statements that correctly describe the capabilities of the Distribution Constructor.

- A. ISO images for use with the Automated Installer (AI) can be created.
- B. Bootable USB images can be created for SPARC and x86 architectures.
- C. A single installation server can be used to create ISO images for SPARC and x86 architectures.
- D. Checkpoints can be used to pause the build, allowing scripts to run that modify the resulting ISO Image.

E. A single Installation server can be used to create ISO images for Solaris 10 and Solaris11 operating systems.

Answer: AD

Explanation:

A: You can use the distribution constructor to create the following types of Oracle Solaris images:

* (A) x86 or SPARC ISO Image for Automated Installations

* Oracle Solaris x86 live CD image

* x86 or SPARC Oracle Solaris text installer image

* x86 Oracle Solaris Virtual Machine

Note: You can use the distribution constructor to build custom Oracle Solaris images. Then, you can use the images to install the Oracle Solaris software on individual systems or multiple systems. You can, also, use the distribution constructor to create Virtual Machine (VM) images that run the Oracle Solaris operating system.

D: Checkpointing Options

You can use the options provided in the `distro_const` command to stop and restart the build process at various stages in the image-generation process, in order to check and debug the image that is being built. This process of stopping and restarting during the build process is called checkpointing.

NEW QUESTION 217

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

1Z0-821 Practice Exam Features:

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

100% Actual & Verified — Instant Download, Please Click
[Order The 1Z0-821 Practice Test Here](#)