



Oracle

Exam Questions 1Z0-821

Oracle Solaris 11 System Administrator

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

The storage pool configuration on your server is:

```
pool1          200K      3.91G   31K    /pool1
pool1/data     31K        3.91G   31K    /pool1/data
pool1          ONLINE      0       0       0
c4t0d0         ONLINE      0       0       0
```

You back up the /pool1/data file system, creating a snapshot and copying that snapshot to tape (/dev/rmt/0). You perform a full backup on Sunday night and incremental backups on Monday through Saturday night at 11:00 pm. Each incremental backup will copy only the data that has been modified since the Sunday backup was started.

On Thursday, at 10:00 am, you had a disk failure. You replaced the disk drive (c4t0d0). You created pool (pool1) on that disk.

Which option would you select to restore the data in the /pool1/data file system?

- A. zfs create pool1/dataLoad the Monday tape and enter:zfs recv pool1/data </dev/rmt/0Load the Wednesday tape and enter:zfs recv -F pool1/data < /dev/rmt/0
- B. Load the Sunday tape and restore the Sunday snapshot:zfs recv pool1/data </dev/rmt/0zfs rollback pool1/data@monLoad the Wednesday tape and restore the Wednesday snapshot:zfs recv -i pool1/data < /dev/rmt/0zfs rollback pool1/data@wed
- C. zfs create pool1/dataLoad the Wednesday tape and enter:zfs recv -F pool1/data </dev/rmt/0
- D. Load the Sunday tape and enter:zfs recv pool1/data < /dev/rmt/0Load the Wednesday tape and enter:* commands missing*

Answer: D

Explanation:

First the full backup must be restored. This would be the Sunday backup.

Then the last incremental backup must be restored. This would be the Wednesday backup. Before restoring the Wednesday incremental file system snapshot, the most recent snapshot must first be rolled back.

By exclusion D) would be best answer even though it is incomplete.

NEW QUESTION 2

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

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

Answer: B

Explanation:

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

NEW QUESTION 3

You are installing Oracle Solaris 11 on a SPARC-based system by using the Test Installer. Which three statements are true?

- A. The ROOT user will always be configured as a role.
- B. The root filesystem will always be deployed on ZFS.
- C. The root filesystem will always be located on a local disk.
- D. The network can be configured using DHCP.
- E. The set of packages that will be installed are server based.
- F. You must always create one regular user when installing the system.

Answer: BDE

NEW QUESTION 4

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 5

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 6

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 7

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

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

Answer: BCD FHI

Explanation:

How to Create an iSCSI LUN

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

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

Example:

```
target# stmfadm create-lu /dev/zvol/rdisk/sanpool/vol1
```

```
Logical unit created: 600144F0B5418B0000004DDAC7C10001
```

4. Confirm that the LUN has been created.

Example

```
target# stmfadm list-lu
```

```
LU Name: 600144F0B5418B0000004DDAC7C10001
```

(F) 5. Add the LUN view.

This command makes the LUN accessible to all systems.

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

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

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

(H) 1. Enable the iSCSI target service.

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

(I) 2. Create the iSCSI target.

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

NEW QUESTION 8

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

What is true?

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

Answer: C

Explanation:

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

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

Permission 7 full

6 read and write

5 read and execute 4 read only

3 write and execute 2 write only

1 execute only

0 none

0 --- no permission 1 --x execute

2 -w- write

3 -wx write and execute 4 r-- read

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

7 rwx read, write and execut

Solaris: Solaris Advanced User's Guide

NEW QUESTION 9

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 10

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

bash: stmfadm: command not found

Which option describes the solution to the problem?

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

Answer: C

Explanation:

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

STMF is controlled by stmfadm, and stmfadm is the majority of the commands you will be using to administer COMSTAR (COMmon Multiprotocl Scsi TARget).

Install the packages you need for COMSTAR with iSCSI and reboot:

```
# pfexec pkg install storage-server
```

```
# pfexec pkg install SUNWiscsit
```

```
# shutdown -y -i6 -g0
```

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

The COMSTAR target mode framework runs as the stmf service. By default, the service is disabled. You must enable the service to use COMSTAR functionality.

You can identify the service with the svcs command. If you have not rebooted the server since installing the group/feature/storage-server package, the service might not be enabled correctly.

NEW QUESTION 10

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 15

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 20

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 21

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

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

Answer: BDFGI

Explanation:

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

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

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

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

G: Install the AI tool set.

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

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

```
-d imagepath
```

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

NEW QUESTION 23

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

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

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

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

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

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

Answer: C

NEW QUESTION 25

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 26

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 29

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

Which method allows you to meet both requirements?

1. Copy live data from /production_ufs to /production_zfs while /production_ufs is in use.
 2. When the copy is complete, /production_zfs will contain an up-to-date copy of /production_ufs
- A. Create a snapshot of the UFS file system
 - B. Create the new ZFS file system
 - C. Use cpio to copy data from the snapshot to the new ZFS file system.
 - D. Create a new Boot Environment
 - E. Create the ZFS file system
 - F. Use lucreate -m to copy data from the Current UFS file system to the new ZFS file system.
 - G. Mirror the existing UFS file system by using SVM. After both submissions are in sync, migrate one of the submissions to a ZFS file system by using Live Upgrade.
 - H. Create the new ZFS file system by using zfs create import to import data from the existing UFS file system into the new ZFS file system

I. Create the new zfs file system by using the zfs create -o shadow.

Answer: E

Explanation:

Migrating Data With ZFS Shadow Migration

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

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

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

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

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

* Create an empty ZFS file system.

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

For example:

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

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

NEW QUESTION 34

_____ 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:

SVCCONFIGD 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 39

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 40

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

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

Answer: DE

Explanation:

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

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

Examples of NVRAM variables:

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

diag-device net The diagnostic startup source device.

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

NEW QUESTION 41

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 45

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 48

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 51

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 56

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 59

User jack logs in to host solar in and issues the following command:

```
jack@solaris:~$ ls .ssh
```

id_dsa id_dsa.pub id_rsa id_rsa.pub known_hosts authorized_keys Which two are true?

- A. The id_rsa file contains the private key for rhosts-based host authentication.
- B. The id_dsa.pub file contains the Digital Signature Algorithm public key for the user jack.
- C. The id_rsa.pub file contains the Rivest Shamir Adelman public key for the host solaris.
- D. The authorized_keys file contains the private keys of remote users authorized to access jack's account on solaris.
- E. The known_hosts file contains the verified public keys of remote hosts known to be trusted.

Answer: AE

Explanation:

A: You will see two files starting with id_rsa. id_rsa is the private key and id_rsa.pub is public key.

E: The .ssh/known_hosts file

In order to use public-key secure connection with other hosts (ssh, scp, sftp) there is a special directory, ~/.ssh/, where passphrases and public keys are stored. Normally you wouldn't need to know the gory details, but from time to time a host will change its public key and then you have difficulty using ssh or scp with that host, and have to edit a file named known_hosts.

If you try to ssh to another computer, but get an error message that warns about a changed or incorrect public key, then it is probably just a case of that host changing its public key. (It is possible, though usually not the case, that malicious hacking is involved.) Unless you actually suspect hacker involvement, you can edit the file ~/.ssh/known_hosts using your usual text editor (vi, emacs, nedit, or pico) and delete any line with the name of that host.

Then when you try to ssh that host again, it will be like the first time ever; ssh will ask you if you want to accept a new public key, you type the whole word yes, and everything will proceed normally from there.

Here is what a typical ~/.ssh/known_hosts file might contain. Note that newton is represented on two different lines:

```
newton 1024 35
153438062610297067329638677441205712613292203533062535600064224677647442
245028855505387934431717435134842994423656065076260604296084868001730665
553662299156116414854701274715680961503198280525759778667306417179500370
189017139564144825610347509023078143132936185076849630461827976942220442
313116255293297021841
ucsub 1024 37
132170811640421742212085598383135714069016332111955003414250071326834884
018721183646445780180633494496866895830879394309011412231102757022090299
732775466435482517698989962531081214859205054227533597152962802400251809
883548442498002326460312850336779152617243800769119880843882425555806081
435017335194477605333
simpson 1024 41
840896920592494584403453622735282634536002054701576247765078766974814128
393752943151071629834843909016027026612791643752972116459602750267266908
365259665072736159491719667576217171370458928680504368847255632477925660
234893185547218857655484574619075125368470792976275806263534208879722192
77539015703446529603
newton, 128.138.249.8 ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEA0d7Aoure0toNJ+YMYi61QP2ka8m5x5ZQIT7obP8C
K3eropfqsmPPY6uiyIh9vpiFX2r1LHcbx139+vG6HOtVvuS8+lfMDtawm3WQvRuOopz3vVy
5GtMwtaOgehsXoT930Ryev1bH5myPtWKlipITsOd2sX9k3tvjrmme4KCGGss=
```

NEW QUESTION 60

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:uid=0
- B. To the file /etc/security/auth_attr, add the line: Network Management: solaris:cmd:RO::/usr/sbin/cryptoadm:uid=0
- C. To the file /etc/security/exec_attr.d/local-entries, add the line: Network Management: solaris:cmd:RO::/usr/sbin/cryptoadm:uid=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 65

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 66

View the Exhibit to inspect the boot environment Information displayed within a non global zone on your system.

BE/Dataset/Snapshot	Active	Mountpoint	Space	Policy	Created
solaris	NR	/	367.97M	static	2011-11-28 11:09
rpool/ROOT/solaris	-	-	26.16M	static	2011-11-28 11:09
rpool/ROOT/solaris/var	-	-	69.0K	static	2011-11-28 13:49
rpool/ROOT/solaris/var@2011-11-28-18:49:38	-	-	0	static	2011-11-28 14:09
rpool/ROOT/solaris/var@2011-11-28-19:09:23	-	-	975.0K	static	2011-11-28 12:29
rpool/ROOT/solaris/var@install	-	-	70.0K	static	2011-11-28 13:49
rpool/ROOT/solaris@2011-11-28-18:49:38	-	-	0	static	2011-11-28 14:09
rpool/ROOT/solaris@2011-11-28-19:09:23	-	-	929.5K	static	2011-11-28 12:29
rpool/ROOT/solaris@install	IR	-	2.0K	static	2011-11-28 13:49
solaris-1	-	-	1.0K	static	2011-11-28 13:49
rpool/ROOT/solaris-1	-	-	-	-	-
rpool/ROOT/solaris-1/var	-	-	57.0K	static	2011-11-28 14:09
z1BE	-	-	1.0K	static	2011-11-28 14:09
rpool/ROOT/z1BE	-	-	-	-	-
rpool/ROOT/z1BE/var	-	-	-	-	-

Which two options describe the solaris-1 boot environment?

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

Answer: AE

Explanation:

A: The - of the Active Column indicates that this boot environment is inactive, and hence not bootable.
 Note: The values for the Active column are as follows: R – Active on reboot.
 N – Active now.
 NR – Active now and active on reboot. "-" – Inactive.
 "!" – Unbootable boot environments in a non-global zone are represented by an exclamation point.
http://docs.oracle.com/cd/E23824_01/html/E21801/unbootable.html#scrolltoc

NEW QUESTION 69

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 72

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 76

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 78

You want to install the openldap software package to a now boot environment for testing before introducing the now software package to the production environment. What option describes the correct procedure to:

- 1) create a new BE named nowBE
- 2) install the software to that new BE only

- A. pkg install --newBE openldap
- B. pkg install --be-nama newBE openldap

- C. beadm create newBEbeadm mount newBE /mntpkg -R /mnt update openldap
- D. beadm create newBEbeadm activate newBEpkg install openldap

Answer: D

Explanation:

If you want to create a backup of an existing boot environment, for example, prior to modifying the original boot environment, you can use the beadm command to create and mount a new boot environment that is a clone of your active boot environment. This clone is listed as an alternate boot environment in the GRUB menu for x86 systems or in the boot menu for SPARC systems.

When you clone a boot environment by using the beadm create command, all supported zones in that boot environment are copied into the new boot environment.

How to Create a Boot Environment

1. Become the root role.
2. Create the boot environment.

```
# beadm create BeName
```

BeName is a variable for the name of the new boot environment. This new boot environment is inactive.

3. (Optional) Use the beadm mount command to mount the new boot environment.

```
# beadm mount BeName mount-point
```

Note: If the directory for the mount point does not exist, the beadm utility creates the directory, then mounts the boot environment on that directory.

If the boot environment is already mounted, the beadm mount command fails and does not remount the boot environment at the newly specified location.

4. (Optional) Activate the boot environment.

```
# beadm activate BeName
```

BeName is a variable for the name of the boot environment to be activated.

On reboot, the newly active boot environment is displayed as the default selection in the x86 GRUB menu or the SPARC boot menu.

NEW QUESTION 81

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 84

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 88

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 92

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 94

dbzone is currently running on your server.

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

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

Answer: DE

Explanation:

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

E: Use: zlogin zone shutdown

to cleanly shutdown the zone by running the shutdown scripts.

Use this procedure to cleanly shut down a zone.

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

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

NEW QUESTION 99

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

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

Answer: B

Explanation:

Boot

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

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

where device-specifier and arguments are optional.

NEW QUESTION 100

Consider the following commands:

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

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

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

Answer: B

NEW QUESTION 104

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 106

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 107

Which statement is correct about shutdown and init commands?

- A. shutdown broadcasts one or more periodic shutdown warning messages to all logged-in users whereas init issues none.
- B. The shutdown command performs a clean shutdown of all services whereas init does not.
- C. The shutdown command brings the system to the single-user milestone by default
- D. The init command must be used to shut the system down to run level 0.
- E. The shutdown command accepts SMF milestones, init stages, or run levels as arguments whereas init accepts only init stages or run levels as arguments.

Answer: A

NEW QUESTION 109

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 111

This icon is displayed on the desktop of a laptop computer, which is running Oracle Solaris 11.



Which two statements describe the Information conveyed by this Icon?

- A. NWAM is disabled.
- B. NWAM is automatically configuring the network.
- C. The wireless network card is manually configured and operational.
- D. The wireless network card is manually configured but not operational.
- E. The wireless network card is automatically configured and operational.
- F. The wireless network card is automatically configured but not operational.

Answer: BC

Explanation:

B: The Network Status notification icon is only displayed on the desktop if you are using NWAM to automatically configure your network.

C: All online (Wireless)

Indicates all manually enabled connections in the enabled network profile are online and that the required number of connections in the enabled profile group (if such a group exists) are online. The required number is the same as those described for the All online (Wired) status.

Note that at least one online connection is wireless.

NEW QUESTION 112

Which three statements accurately describe the Automated Installation (AI) client?

- A. If the AI client does not match any criteria to use a custom manifest or script, the default manifest is used.
- B. If the AI client does not match any criteria to use a custom manifest or script, the automated installation aborts.
- C. Any manifest or script in a service can be designated to be the default for that service.
- D. Only the default.xml file is used as the default AT client manifest.
- E. If a client system does not use any SC profile, then an interactive tool opens on that client at first boot after that client installation to complete the configuration of that client.
- F. If a client system does not use any SC profile, then the install server will use the default SC profile.

Answer: ACE

Explanation:

Each client uses one and only one AI manifest to complete its installation. The AI manifest is selected for a client according to the following algorithm:

* If no custom AI manifests are defined for this install service, the default AI manifest is used. The default AI manifest is not associated with any client criteria etc.

Each client can use any number of system configuration profiles. If a client system does not use any configuration profile, then an interactive tool opens on that client at first boot after that client installation to complete the configuration of that client.

NEW QUESTION 113

You are executing this command in the default shell: `sleep 5000 &`
 The system displays a number. This value is .

- A. the priority of the `/usr/bin/sleep` process
- B. the process ID of the `/usr/bin/sleep` process
- C. the process ID of the shell spawned to execute `/usr/bin/sleep`
- D. the process group ID that includes the `/usr/bin/sleep` process
- E. the amount of memory allocated to the `/usr/bin/sleep` process
- F. the current number of instances of the `/usr/bin/sleep` process

Answer: C

Explanation:

If a command is terminated by the control operator '&', the shell executes the command asynchronously in a subshell. This is known as executing the command in the background. The shell does not wait for the command to finish, and the return status is 0 (true).

NEW QUESTION 116

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

The zone has never been booted. Which three options correctly describe this zone?

- A. It is a sparse root zone.
- B. It is a whole root zone.
- C. It is an immutable zone.
- D. It is a native zone.
- E. The zone shares the network interface with the host.
- F. The zone uses a virtual network interface.
- G. The hostid is the same as the global zone.
- H. The IP address of the zone is 10.0.2.18.

Answer: CEG

Explanation:

C: Immutable Zones provide read-only file system profiles for solaris non-global zones. Note that ip-type: exclusive: Starting with OpenSolaris build 37 and Oracle Solaris 10 8/07, a default zone can be configured as an "exclusive-IP zone" which gives it exclusive access to the NIC(s) that the zone has been assigned. Applications in such a zone can communicate directly with the NIC(s) available to the zone.

Note on zones:

After installing Oracle Solaris on a system, but before creating any zones, all processes run in the global zone. After you create a zone, it has processes that are associated with that zone and no other zone. Any process created by a process in a non-global zone is also associated with that non-global zone.

Any zone which is not the global zone is called a non-global zone. Most people call non- global zones simply "zones." Some people call them "local zones" but this is discouraged.

The default native zone file system model on Oracle Solaris 10 is called "sparse-root." This model emphasizes efficiency and security at the cost of some configuration flexibility. Sparse-root zones optimize physical memory and disk space usage by sharing some directories, like /usr and /lib. Sparse-root zones have their own private file areas for directories like /etc and /var. Whole-root zones increase configuration flexibility but increase resource usage. They do not use shared file systems for /usr, /lib, and a few others.

There is no supported way to convert an existing sparse-root zone to a whole-root zone. Creating a new zone is required.

NEW QUESTION 126

The /etc/hosts file can be best described as .

- A. a local database of host names for rlogin, rsh, and rep
- B. the configuration file for the host name of the system
- C. a local database of information for the uname command
- D. the configuration file for the Domain Name Service (DNS)
- E. a local database of host names and their associated IP addresses

Answer: E

Explanation:

As your machine gets started, it will need to know the mapping of some hostnames to IP addresses before DNS can be referenced. This mapping is kept in the /etc/hosts file. In the absence of a name server, any network program on your system consults this file to determine the IP address that corresponds to a host name.

NEW QUESTION 128

You wish to troubleshoot some issues that you are having on the system. You want to monitor the /var/adm/messages file in real time. Which command would you use to do this?

- A. head
- B. tail
- C. cat
- D. file
- E. test

Answer: B

Explanation:

tail is a program on Unix and Unix-like systems used to display the last few lines of a text file or piped data.

By default, tail will print the last 10 lines of its input to the standard output. With command line options the number of lines printed and the printing units (lines, blocks or bytes) may be changed. The following example shows the last 20 lines of filename:

tail -n 20 filename

NEW QUESTION 133

zone1 is a non-global zone that has been configured and installed.

zone1 was taken down for maintenance, and the following command was run: zoneadm -z zone1 mark incomplete

The following information is displayed when listing the zones on your system:

ID	NAME	STATUS	PATH	BRAND	IP
0	global	running	/	solaris	shared
-	dbzone	installed	/export/dbzone	solaris	excl
-	zone1	incomplete	/zone/zone1	solaris10	excl

Which task needs to be performed before you can boot zone1?

- A. The zone needs to be installed.
- B. The zone needs to be brought to the ready state.
- C. The zone needs to be uninstalled and reinstalled.
- D. The zone needs to be brought to the complete state.

Answer: C

Explanation:

If administrative changes on the system have rendered a zone unusable or inconsistent, it is possible to change the state of an installed zone to incomplete. Marking a zone incomplete is irreversible. The only action that can be taken on a zone marked incomplete is to uninstall the zone and return it to the configured state.

NEW QUESTION 137

You have a process called bigscript, and you need to know the PID number for this process. Which command will provide that information?

- A. pkill bigscript
- B. ps bigscript
- C. pgrep bigscript
- D. prstat bigscript

Answer: C

Explanation:

Pgrep takes a process name and return a PID.

Note: pgrep looks through the currently running processes and lists the process IDs which matches the selection criteria to stdout. All the criteria have to match. For example, pgrep - u root sshd will only list the processes called sshd AND owned by root.

Incorrec answers:

ps bigscript: You can't pass a name to ps, it interprets it as arguments.

NEW QUESTION 141

Which two options accurately describe the network characteristics of a zone?

- A. DHCP address assignment cannot be configured in a shared IP zone.
- B. Shared IP is the default type of network configuration.
- C. Exclusive IP is the default type of network configuration.
- D. By default, all IP addresses, netmasks, and routes are set by the global zone and cannot be altered in a non global zone.
- E. IPMP cannot be managed within the non-global zone.
- F. Commands such as snoop and dladm cannot be used on datalinks that are in use by a running zone.

Answer: AB

Explanation:

A: Non-global zones can not utilize DHCP (neither client nor server).

B (not C): By default, non-global zones will be configured with a shared IP functionality. What this means is that IP layer configuration and state is shared between the zone you're creating and the global zone. This usually implies both zones being on the same IP subnet for each given NIC.

Note: A zone is a virtual operating system abstraction that provides a protected environment in which applications run. The applications are protected from each other to provide software fault isolation. To ease the labor of managing multiple applications and their environments, they co-exist within one operating system instance, and are usually managed as one entity.

The original operating environment, before any zones are created, is also called the "global zone" to distinguish it from non-global zones, The global zone is the operating system instance.

Incorrect Answer

E: Exclusive-IP zones can use IPMP. IPMP is configured the same way in an exclusive-IP zone as it is on a system not using zones.

For shared-IP zones, IPMP can be configured in the global zone. F: Full IP-level functionality is available in an exclusive-IP zone. An exclusive-IP zone has its own IP-related state.

An exclusive-IP zone is assigned its own set of data-links using the zonecfg command. The zone is given a data-link name such as xge0, e1000g1, or bge32001, using the physical property of the net resource. The address property of the net resource is not set.

Note that the assigned data-link enables the snoop command to be used.

The dladm command can be used with the show-linkprop subcommand to show the assignment of data-links to running exclusive-IP zones.

NEW QUESTION 146

.....

Relate Links

100% Pass Your 1Z0-821 Exam with ExamBible Prep Materials

<https://www.exambible.com/1Z0-821-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>