



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

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.

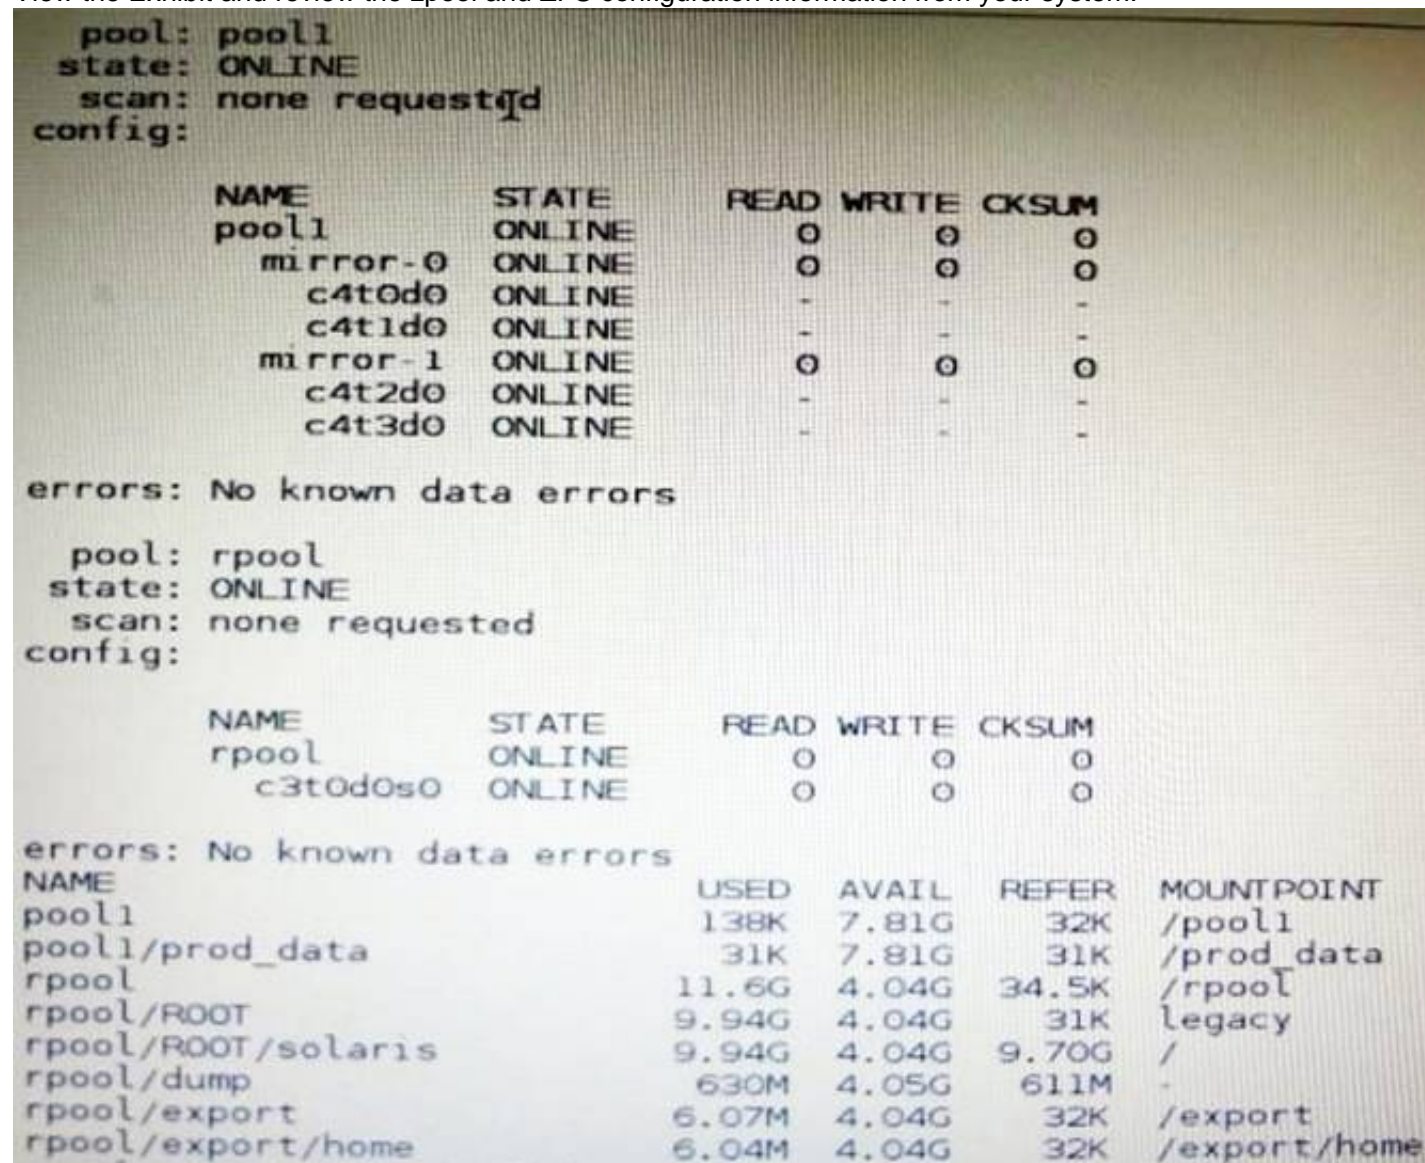
beadm list

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.



```
pool: pool1
state: ONLINE
scan: none requested
config:

    NAME                STATE                READ WRITE CKSUM
    pool1                ONLINE              0     0     0
      mirror-0           ONLINE              0     0     0
        c4t0d0           ONLINE              -     -     -
        c4t1d0           ONLINE              -     -     -
      mirror-1           ONLINE              0     0     0
        c4t2d0           ONLINE              -     -     -
        c4t3d0           ONLINE              -     -     -

errors: No known data errors

pool: rpool
state: ONLINE
scan: none requested
config:

    NAME                STATE                READ WRITE CKSUM
    rpool                ONLINE              0     0     0
      c3t0d0s0           ONLINE              0     0     0

errors: No known data errors

NAME                                USED    AVAIL    REFER    MOUNTPOINT
pool1                               138K    7.81G    32K      /pool1
pool1/prod_data                     31K    7.81G    31K      /prod_data
rpool                               11.6G    4.04G    34.5K    /rpool
rpool/ROOT                          9.94G    4.04G    31K      legacy
rpool/ROOT/solaris                  9.94G    4.04G    9.70G    /
rpool/dump                          630M    4.05G    611M    -
rpool/export                       6.07M    4.04G    32K      /export
rpool/export/home                   6.04M    4.04G    32K      /export/home
```

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

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 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

You are troubleshooting network throughput on your server.

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

The correct command is .

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

Answer: B

Explanation:

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

NEW QUESTION 6

You are going to create live zones on you server. Disk space is critical on this server so you need to reduce the amount of disk space required for these zones.

Much of the data required for each of these zones is identical, so you want to eliminate the duplicate copies of data and store only data that is unique to each zone.

Which two options provide a solution for eliminating the duplicate copies of data that is common between all of these zones?

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

- E. Put each zone in a separate ZFS file system within the same zpoo
- F. Set the dedup property to on for each ZFS file system.

Answer: DE

Explanation:

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

NEW QUESTION 7

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 8

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 Multiprotocol Scsi TARget). Install the packages you need for COMSTAR with iSCSI and reboot:

```
# pfexec pkg install storage-server
# pfexec pkg install SUNWiscsit
# shutdown -y -i6 -g0
```

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

The COMSTAR target mode framework runs as the stmf service. By default, the service is disabled. You must enable the service to use COMSTAR functionality. You can identify the service with the svcs command. If you have not rebooted the server since installing the group/feature/storage-server package, the service might not be enabled correctly.

NEW QUESTION 9

User jack on host solaris attempts to use ssh to log in to host oracle and receives this message:
jack@solaris:~\$ ssh oracle
ssh: connect to host oracle port 22: connection refused What is the problem?

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

Answer: G

Explanation:

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

NEW QUESTION 10

You have Solaris 11 system with a host name of sysA and it uses LDAP as a naming service. You have created a flash archive of sysA and you want to migrate this system to an Oracle Solaris11 server, Solaris10 branded zone.

The zone Status on the Oracle Solaris 11 server is:

- zone10 incomplete/zone/zone1solaris10exc1

Select the option that will force the non-global zone to prompt you for a host name and name service the first time it is booted.

- A. Use zonecfg to change the zonename before booting the system for the first time
- B. Use the -u option with the zoneadm -z zone10 attach command.
- C. Use the -u option with the zoneadm -z zone10 install command.
- D. Remove the sysidcfg file from the <zonepath>/root directory before booting the non- global zone.

Answer: C

Explanation:

Oracle Solaris 10 branded zones – Oracle Solaris 10 Zones provide an Oracle Solaris 10

environment on Oracle Solaris 11. You can migrate an Oracle Solaris 10 system or zone to a solaris10 zone on an Oracle Solaris 11 system in the following ways:

* Create a zone archive and use the archive to create an s10zone on the Oracle Solaris 11 system.

This option applies in the current scenario.

Example of command to Install the Oracle Solaris 10 non-global zone. s11sysB# zoneadm -z s10zone install -u -a /pond/s10archive/s10.flar

* Detach the zone from the Oracle Solaris 10 system and attach the zone on the Oracle Solaris 11 zone. The zone is halted and detached from its current host.

The zonepath is moved to the target host, where it is attached.

Note:

install [-x nodataset] [brand-specific options] A subcommand of the zoneadm.

Install the specified zone on the system. This subcommand automatically attempts to verify first. It refuses to install if the verify step fails.

-u uuid-match

Unique identifier for a zone, as assigned by libuuid(3LIB). If this option is present and the argument is a non-empty string, then the zone matching the UUID is selected instead of the one named by the -z option, if such a zone is present.

NEW QUESTION 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 14

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

pgrep userprogram l5317

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

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

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

Answer: B

Explanation:

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

To terminate a process:

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

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

NEW QUESTION 15

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

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

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

pkg/inst_root property and the pkg/readonly property to /export/IPS/reposet the preferred pkg/inst_root property by using pkg set-publisher - G
http://serverA.example.com/ \- g http://pkg.oracle.com/solaris.com/release/- p solaris

Answer: A

NEW QUESTION 20

Given the following output of the zpool status command:

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

Identify the correct statement regarding pool1's configuration.

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

Answer: B

NEW QUESTION 25

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 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 /ptoduction_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 syste
- B. Create the new ZFS file syste
- C. Use cpio to copy data from the snapshot to the new ZFS file system.
- D. Create a new Boot Environmen
- E. Create the ZFS file syste
- 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:

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 38

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

A. Use the V2V process to migrate an existing Solaris 8 or 9 non-global zone from a Solaris 10 system to a solaris10 branded zone.

B. Use the V2V process to migrate an existing Solaris 10 non-global whole root zone from a Solaris 10 system to a solaris10 branded whole root zone.

C. Install a solaris10 branded zone directly from the Oracle Solaris 10 media.

D. Migrate an existing 64-bit Solaris 10 system to a solaris10 branded non-global zone using the P2V process.

E. Use the V2V process to migrate an existing Solaris 10 non-global sparse root zone from a Solaris 10 system to a solaris10 branded sparse root zone.

Answer: BC

Explanation:

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

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

C: How to Install the solaris10 Branded Zone

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

NEW QUESTION 39

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 41

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

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

You then enter: `ipadm create-ip net3`

What is the output?

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

Answer: B

Explanation:

According to the exhibit the interface already exists.

The command ipadm create-ip net3 is supposed to create a new interface net3.

NEW QUESTION 44

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 48

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

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

export genius

In the .bashrc us this line: genius=kernighan

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

export genius

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

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

Answer: C

NEW QUESTION 49

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 54

You created a new zpool. Now you need to migrate the existing ZFS file system from pool1/prod to pool2/prod.

You have these requirements:

- 1. Users must have access to the data during the migration, so you cannot shutdown the file system while the migration takes place.
- 2. Because you want to copy the data as quickly as possible, you need to increase the server resources devoted to the ZFS migration.

Which method would you use to modify the ZFS shadow migration daemon defaults to increase the concurrency and overall speed of migration?

- A. Svccfg - s filesystem/shadowd:defaultsetprop config_params/shadow_threads=integer: 16endsvcadm refresh filesystem/shadowd: default
- B. Specify the -b <blocksize> option with the zfs create command and increase the value of<blocksize>

- C. Use the -o -volblocksize=<blocksize>option with the zfs create command and increase the value of the default <blocksize>.
D. Svccfg -s filesystem/zfs: defaultsetprop config_params/shadow_threads = integer: 16endsvcadm refresh filesystem/zfs:default

Answer: A

Explanation:

shadowd is a daemon that provides background worker threads to migrate data for a shadow migration. A shadow migration gradually moves data from a source file system into a new “shadow” file system. Users can access and change their data within the shadow file system while migration is occurring.

The shadowd service is managed by the service management facility, smf(5).

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). The service's status can be queried using the svcs(1) command.

The svccfg(1M) command can be used to manage the following parameter related to shadowd:

config_params/shadow_threads

Note: Oracle Solaris 11: In this release, you can migrate data from an old file system to a new file system while simultaneously allowing access and modification of the new file system during the migration process.

Setting the shadow property on a new ZFS file system triggers the migration of the older data. The shadow property can be set to migrate data from the local system or a remote system with either of the following values:

file:///path nfs://host:path

NEW QUESTION 59

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

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

Answer: B

Explanation:

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

Compared to TCP, UDP is a simpler message-based connectionless protocol.

NEW QUESTION 63

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 66

The following image properties are displayed on your system:

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

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

- A. All package operations are performed in a new BE set as active on the next boot.
B. Do not create a new B

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

Answer: DF

Explanation:

Image properties described below.

* be-policy

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

/ default

Apply the default BE creation policy: create-backup.

/ always-new (D, F)

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

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

NEW QUESTION 70

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  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 72

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

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

Answer: D

Explanation:

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

```
# ipadm show-addr
```

```
ADDROBJ TYPE STATE ADDR
```

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

NEW QUESTION 77

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 81

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

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

Answer: B

Explanation:

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

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

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

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

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

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

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

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

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

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

NEW QUESTION 83

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.

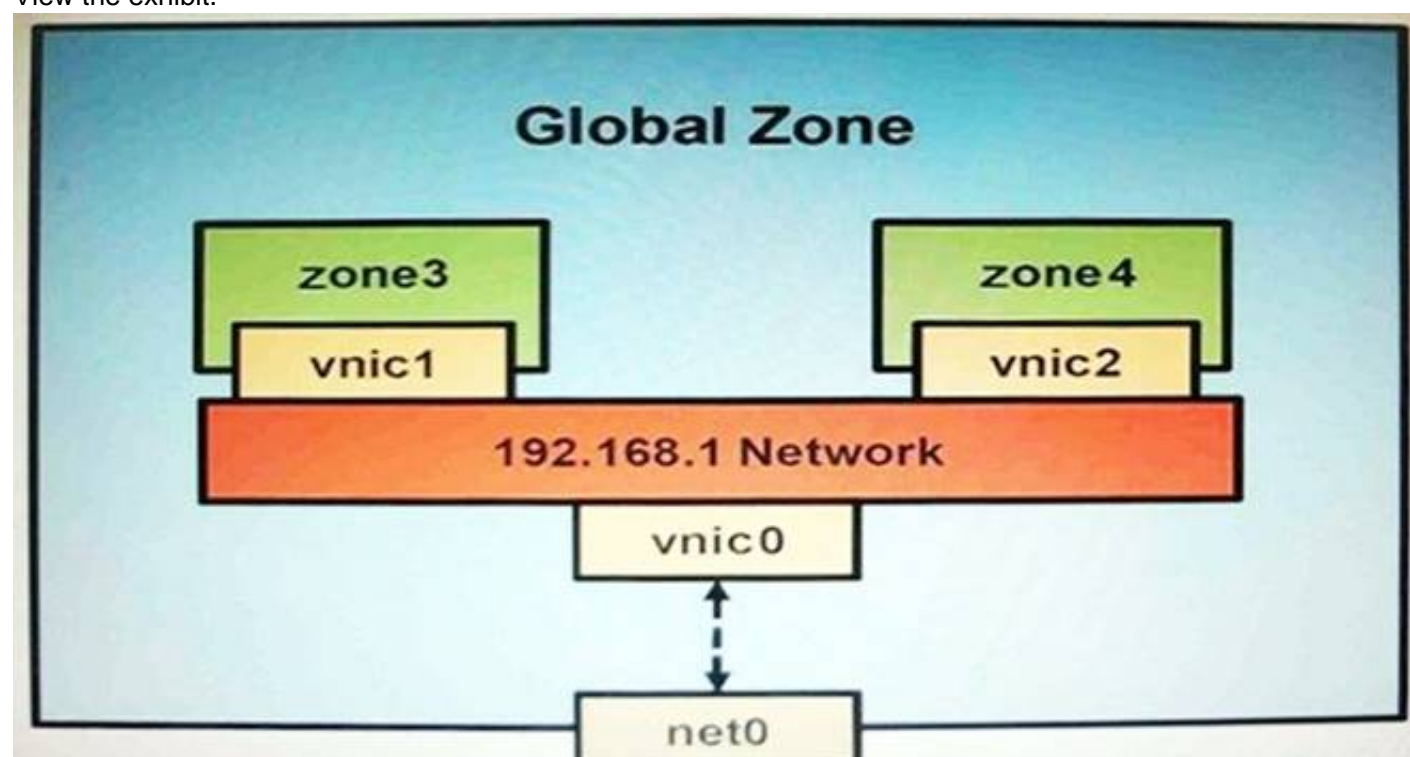
1. Install the COMSTAR storage server software. target# pkg install storage-server

Etc.

NEW QUESTION 84

You have been asked to troubleshoot the initial configuration of a virtual network connecting two local zones with the outside world.

View the exhibit.



The command

dladm create-vnic -1 vswitch192.168.1 vnic1 fails with the error

dladm: invalid link name 'vswitch192.168.1' What is the reason for this error?

- A. The name vswitch192.168.1 is not legal.
- B. The zone must be specified withdladm create-vnic -z zone3 vnic1.
- C. The virtual interface must be specified withdladm create-vnic -z zone3 vnic1.
- D. The virtual interface must be created withipadm create-vnic -1 switch192.168.1.

E. The virtual switch must be created first with `dladm create -etherstub vswitch192.168.1`.

Answer: E

Explanation:

There is no data-link named `vswitch192.168`. We need to create an etherstub first.

See Note and example below for details.

Note: Create a VNIC in the system's global zone.

`# dladm create-vnic -l data-link vnic-name`

`data-link` is the name of the interface where the VNIC is to be configured.

`-l link, --link=link`

`link` can be a physical link or an etherstub.

`vnic-name` is the name that you want to give the VNIC.

For example, to create a VNIC named `vnic0` on interface `e1000g0`, you would type the following:

`# dladm create-vnic -l e1000g0 vnic0`

Example: Creating a Virtual Network Without a Physical NIC First, create an etherstub with name `stub1`:

`# dladm create-etherstub stub1`

Create two VNICs with names `hello0` and `test1` on the etherstub. This operation implicitly creates a virtual switch connecting `hello0` and `test1`.

`# dladm create-vnic -l stub1 hello0`

`# dladm create-vnic -l stub1 test1`

NEW QUESTION 88

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
Ä ufsdump level zero (full) backups

Note:

Branded zones that run an environment different than the OS release on the system

* The `lx` branded zone introduced in the Solaris 10 8/07 release provides a Linux environment for your applications and runs on x86 and x64 machines on the Oracle Solaris 10 OS.

* The `solaris8` and `solaris9` branded zones enable you to migrate an Oracle Solaris 8 or Oracle Solaris 9 system to an Oracle Solaris 8 or Oracle Solaris 9 Container on a host running the Oracle Solaris 10 8/07 Operating System or later Oracle Solaris 10 release.

* The Oracle Solaris 10 Container brand is available in OpenSolaris build 127. These branded zones host Oracle Solaris 10 user environments.

Note: One of the powerful features of Solaris 11 is the ability to run a Solaris 10 environment in a zone. Solaris 10 allows you to run Solaris 8 and 9 environments in zones, but only on SPARC.

NEW QUESTION 92

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 97

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 101

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 104

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 109

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 110

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 113

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 118

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

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

Answer: B

Explanation:

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

```
cat /etc/release
```

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

Example:

```
# cat /etc/release
```

```
Oracle Solaris 10 8/11 s10s_u10wos_17b SPARC
```

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

NEW QUESTION 120

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 123

Your task is to convert a JumpStart sysidcfg file to an Automated Installer (AI) sc_profile.xml file, using js2ai. Select two unsupported items that will require changes.

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

Answer: AD

Explanation:

A: terminal

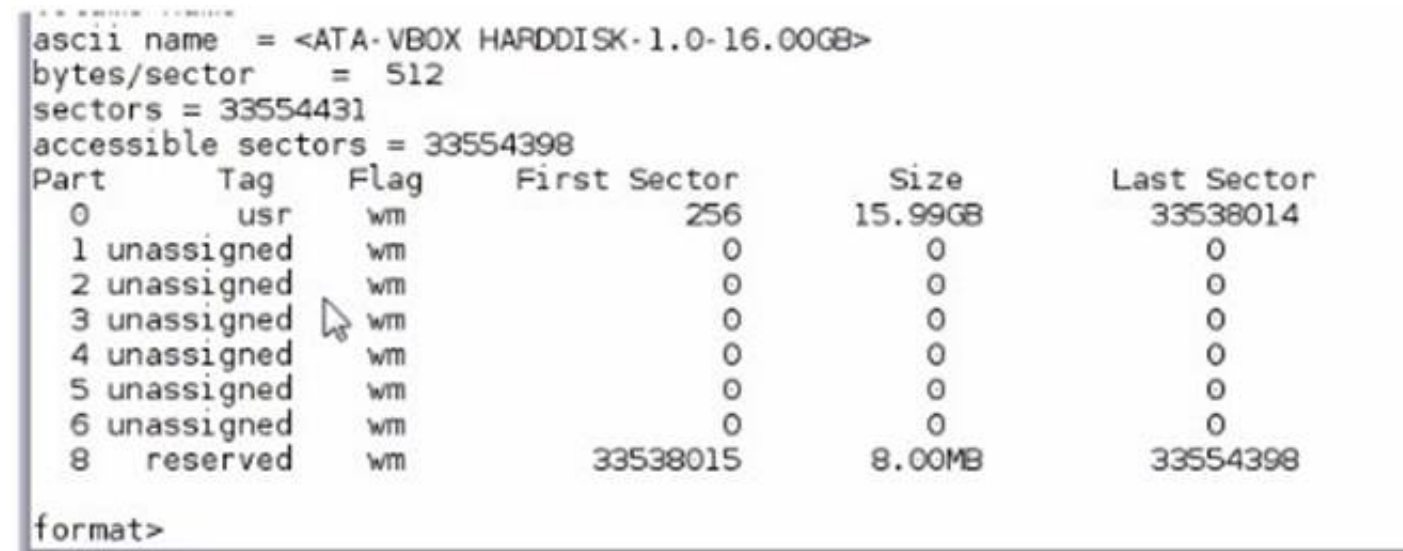
The js2ai tool does not perform any translation. Make sure the terminal type specified in the sysidcfg file is supported in Oracle Solaris 11.

D: system_locale

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

NEW QUESTION 125

View the Exhibit.



The screenshot shows a disk configuration utility. At the top, it displays disk information: `ascii name = <ATA-VBOX HARDDISK-1.0-16.00GB>`, `bytes/sector = 512`, `sectors = 33554431`, and `accessible sectors = 33554398`. Below this is a table with columns: Part, Tag, Flag, First Sector, Size, and Last Sector. The table contains 9 rows of data. The first row (Part 0) is marked 'usr' and 'wm', with a first sector of 256, a size of 15.99GB, and a last sector of 33538014. The next seven rows (Parts 1-7) are marked 'unassigned' and 'wm', with a first sector of 0, a size of 0, and a last sector of 0. The final row (Part 8) is marked 'reserved' and 'wm', with a first sector of 33538015, a size of 8.00MB, and a last sector of 33554398. At the bottom, the prompt `format>` is visible.

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
7	unassigned	wm	0	0	0
8	reserved	wm	33538015	8.00MB	33554398

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 128

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 129

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 133

Review the non-global zone configuration displayed below:

```
zonename: dbzone
zonepath: /export/dbzone
brand: Solaris
autoboot: false
bootargs:
file-mac-profile:
pool:
limitpriv:
scheduling-class:
ip-type: exclusive
hostid:
fs-allowed:
anet:
    linkname: net0
    lower-link: auto
    allowed-address not specified
    configure-allowed-address: true
    defrouter not specified
    allowed-dhcp-cids not specified
    link-protection: mac-nospoof
    mac-address: random
    auto-mac-address: 2:8:20:97:40:20
    mac-prefix not specified
    mac-slot not specified
    vlan-id not specified
    priority not specified
    rxrings not specified
    rxrings not specified
    mtu not specified
    maxlow not specified
    rxfanout not specified
```

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

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

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

Answer: C

Explanation:

Add a memory cap.

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

NEW QUESTION 138

Server A, Server B, and Server C are connected to the same network switch and are on the same network. 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 B, you enter the following command: `dladm show-phys | grep net0`

Response:

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

What is the next logical troubleshooting action?

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

Answer: D

Explanation:

Check the physical connection.

NEW QUESTION 141

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 device set match=/dev/rdisk/c4t0d0` End Boot the zone and log in the zone as root
- B. Create the zpool: `zpool create pool2 c4t0d0` In the non-global zone, root can now create ZFS file system in the pool2 zpool
- C. In the global zone, create the zpool: `global# zpool create pool2 c4t1d0` While creating the non-global zone, make the following entry: `add dataset set name=pool2` End `add fs set dir=pool1 set special=pool1 set type=zfs` pool1 End Boot the zone, log in the zone as root, and create the zfs file system in the pool2 zpool.
- D. In the global zone, create the zpool: `global# zpool create pool2 c4t1d0` While creating the global zone, make the following entry: `add dataset set name=pool2` End Boot 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 c4t1d0` `global# zfs create pool2/data` While creating the non-global zone, make the following entry for each ZFS file system that you want to make available in the zone: `add fs set dir=/dataset special=pool2/dataset type=zfs` End
- F. Create the zpool in the global zone: `global# zpool create pool2 c4t1d0` Boot the non-global zone, log in to the zone as root, and issue this command to delegate ZFS permissions to root: `non-global zone# zfs allow root create , destroy, mount` pool2 Log in to the non-global zone create ZFS file systems in the pool2 zpool.

Answer: C

Explanation:

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

NEW QUESTION 145

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 149

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

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

Answer: A

NEW QUESTION 151

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

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

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

Answer: BD

Explanation:

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

LOG_FROM_REMOTE

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

NEW QUESTION 152

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 156

The su command by default makes an entry into the log file for every su command attempt. The following is a single line from the file:

SU 12/18 23:20 + pts/1 user1-root What does the + sign represent?

- A. unsuccessful attempt
- B. successful attempt
- C. The attempt was from a pseudo terminal, and not the console.
- D. The attempt was from a user that is in the adm group, same as root.
- E. Time zone is not set.

Answer: B

Explanation:

The sulog file, /var/adm/sulog, is a log containing all attempts (whether successful or not) of the su command. An entry is added to the sulog file every time the su command is executed. The fields in sulog are: date, time, successful (+) or unsuccessful (-), port, user executing the su command, and user being switched to. In the preceding example, all su attempts were successful, except for the attempt on 2/23 at 20:51, when user pete unsuccessfully attempted to su to user root. Look for entries where an unauthorized user has used the command inappropriately. The following entry shows a successful (indicated by +) su from user userid to root.

SU 03/31 12:52 + pts/0 <userid>-root

NEW QUESTION 159

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 syste
- B. When the system is booted, execute the zpool clear pool1 command.
- C. Shut the system down, replace disk c3t3d0, and boot the syste
- D. When the system is booted execute the zpool online pool1 command.
- E. Shut the system down, replace disk c3t3d0, and boot the syste
- F. When the system is booted, execute the zpool replace pool1 c3t3d0 command.
- G. Shut the system down, replace disk c3t3d0, and boot the syste
- 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::disk/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 163

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 168

Subnets are created by using .

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

Answer: B

Explanation:

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

The routing prefix of an address is written in a form identical to that of the address itself. This is called the network mask, or netmask, of the address. For example, a specification of the most-significant 18 bits of an IPv4 address, 11111111.11111111.11000000.00000000, is written as 255.255.192.0.

NEW QUESTION 169

You have completed configuring a zone named dbzone on your Solaris 11 server. The configuration is as following:

```
zonename: dbzone
zonepath: /export/dbzone
brand: Solaris
autoboot: false
bootargs:
file-mac-profile:
pool:
limitpriv:
scheduling-class:
ip-type: exclusive
hostid:
fs-allowed:
anet:
    linkname: net0
    lower-link: auto
    allowed-address not specified
    configure-allowed-address: true
    defrouter not specified
    allowed-dhcp-cids not specified
    link-protection: mac-nospoof
    mac-address: random
    mac-prefix not specified
    mac-slot not specified
    vlan-id not specified
    priority not specified
    rxrings not specified
    rxrings not specified
    mtu not specified
    maxlow not specified
    rxfanout not specified
```

The global zone displays the following network information:

ADDROBJ	TYPE	STATE	ADDR
lo0/v4	static	ok	127.0.0.1/8
net0/_b	dhcp	ok	10.0.2.18/24
lo0/v6	static	ok	::1/128
net0/_a	addrconf	ok	fe80::a00:27ff:fe8e:c0d4/10

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 173

You have set up the task.max-lwps resource control on your Solaris 11 system.

Which option describes how to configure the system so that syslogd notifies you when the resources control threshold value for the task.max-lwps resource has been exceeded?

- A. Use the rctldm command to enable the global action on the task.max-lwpa resource control.
- B. Modify the /etc/syslog.conf file to activate system logging of all violations of task.max-lwps and then refresh then svc: /system/system-log:default service.
- C. Activate system logging of all violations of task.max-lwpp in the /etc/rctldm.conf file and then execute the rctldm-u command.
- D. Use the prct1 command to set the logging of all resource control violations at the time the task.max-lwps resource control is being setup.
- E. Use the setrct1 command to set the logging of all resource control violations for the task.max-lwps resource control.

Answer: A

Explanation:

rctldm - display and/or modify global state of system resource controls

The following command activates system logging of all violations of task.max-lwps.

```
# rctldm -e syslog task.max-lwps
```

```
#
```

NEW QUESTION 175

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

ADDBOBJ	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 180

You are asked to troubleshoot networking issues on an unfamiliar system. Select the correct command to display what network devices are installed.

- A. ifconfig -a
- B. dladm show-dev
- C. dladm show-phys
- D. dladm show-ether
- E. netadm show-dev
- F. netadm show-ether

Answer: C

NEW QUESTION 183

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
	mirror-1	ONLINE	0	0	0
	c3t5d0	ONLINE	0	0	0
	c3t6d0	ONLINE	0	0	0

The following partition scheme is used for every disk drive in pool1:

```
ascii name = <ATA-VBOX HARDDISK-1.0-146.00GB>
bytes/sector = 512
sectors = 306184191
accessible sectors = 306184158
```

Part	Tag	Flag	First Sector	Size	Last Sector
0	usr	wm	256	145.99GB	306167774
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	306167775	8.00MB	306184158

Which two are true regarding the ZFS storage pool?

- A. The data on c3t3d0 is duplicated on c3t4d0.
- B. The data is striped across disks c3t3d0 and c3t4d0 and mirrored across vdevs mirror-0 and mirror-1.
- C. The storage pool is 146 GB total size (rounded to the nearest GB).
- D. The storage pool is 584 GB total size (rounded to the nearest GB).
- E. The storage pool is 292 GB total size (rounded to the nearest GB).

Answer: AE

NEW QUESTION 185

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

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

Which two statements are correct?

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

Answer: AC

Explanation:

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

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

NEW QUESTION 189

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

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

Answer: B

Explanation:

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

NEW QUESTION 194

On localSYS, your SPARC based server, you back up the root file system with recursive snapshots of the root pool. The snapshots are stored on a remote NTS file system.

This information describes the remote system where the snapshots are stored:

Remote system name: backupSYS

File system where the snapshots are stored: /backups/localSYS Mounted file system on localSYS: /rpool/snaps

Most recent backup name: rpool-1202

Disk c0t0d0 has failed in your root pool and has been replaced. The disk has already been

parted and labeled and now you need to restore the root file system. Which procedure would you follow to restore the ZFS root file system on localSYS?

- A. boot cdrom -smount -f nfs backup_server:/rpool/snaps /rmtzpool create rpool c0t0d0s0cat /mnt/rpool.1202 | zfs receive -Fdu rpoolzpool set bootfs=rpool/ROOT/solaris rpoolRecreate swap and dump devices.Reinstall the bootblock on c0t0d0.
- B. boot cdrom -smount -f nfs backup_server:/rpool/snaps /mntzpool create rpool c0t0d0s0zfs create -o mountpoint=/ rpool/ROOTcat /mnt/rpool.1011 | zfs receive -Fdu rpoolzpool set bootfs=rpool/ROOT/solaris rpoolRecreate swap and dump devices.Reinstall the bootblock on c0t0d0.
- C. boot cdrom -smount -F nfs backup_server:/rpool/snaps /mntcat /mnt/rpool.1011 | zfs receive -Fdu rpoolzpool set bootfs=rpool/ROOT/solaris rpool c0t0d0s0Reinstall the bootblock on c0t0d0s0
- D. boot cdrom -smount -f nfs backup_server:/rpool/snaps /rmtzpool create rpool c0t0d0s0zfs receive -Fdu /mnt/rpool.1011zpool set bootfs=rpool/ROOT/solaris rpoolReinstall the bootblock on c0t0d0.

Answer: A

Explanation:

How to Recreate a ZFS Root Pool and Restore Root Pool Snapshots In this scenario, assume the following conditions:

- * ZFS root pool cannot be recovered
- * ZFS root pool snapshots are stored on a remote system and are shared over NFS
- * The system is booted from an equivalent Solaris release to the root pool version so that the Solaris release and the pool version match. Otherwise, you will need to add the -o version=version-number property option and value when you recreate the root pool in step 4 below.

All steps below are performed on the local system. 1.

Boot from CD/DVD or the network.

On a SPARC based system, select one of the following boot methods:

ok boot net -s

ok boot cdrom -s

If you don't use -s option, you'll need to exit the installation program.

2.

Mount the remote snapshot dataset. For example:

```
# mount -F nfs remote-system:/rpool/snaps /mnt
```

3.

Recreate the root pool. For example:

```
# zpool create -f -o failmode=continue -R /a -m legacy -o cachefile=/etc/zfs/zpool.cache rpool c1t0d0s0
```

4.

Restore the root pool snapshots.

This step might take some time. For example:

```
# cat /mnt/rpool.0311 | zfs receive -Fdu rpool
```

Using the -u option means that the restored archive is not mounted when the zfs receive operation completes.

5.

Set the bootfs property on the root pool BE. For example:

```
# zpool set bootfs=rpool/ROOT/osalBE rpool 6.
```

Install the boot blocks on the new disk.

On a SPARC based system:

```
# installboot -F zfs /usr/platform/`uname -i`/lib/fs/zfs/bootblk /dev/rdisk/c1t0d0s0
```

NEW QUESTION 196

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 198

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 202

Before booting test zone 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 - testzone - c

Answer: E

Explanation:

The following options are supported:

C

Connects to the zone console. Connects to the zone console.

Note:

After you install a zone, you must log in to the zone to complete its application environment. You might log in to the zone to perform administrative tasks as well.

Unless the -C option is used to connect to the zone console, logging in to a zone using zlogin starts a new task. A task cannot span two zones

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 210

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

You upgraded your server to Oracle Solaris 11 and you imported zpool (pool1) that was created in Solaris 10. You need to create an encrypted ZFS file system in pool1, but first you need to make sure that your server supports ZFS encryption. Which four statements are true for support of ZFS encryption?

- A. The encrypted file system must have been created in Oracle Solaris11. To encrypt a ZFS file system from a previous version of Solaris, upgrade the zpool and create a new encrypted ZFS file system into the encrypted ZFS file system.
- B. If you plan to create an encrypted file system in an existing zpool, the zpool must be upgraded to ZFS version 30.
- C. ZFS encryption is integrated with the ZFS command set and no additional packages need to be installed.
- D. ZFS encryption requires that the ZFS Dataset Encryption package be installed.
- E. If you plan to create an encrypted file system in an existing zpool, the pool must be upgraded to ZFS version 21, minimum.
- F. Encryption is supported at the pool or dataset (file system) level.
- G. Encryption is supported at the pool level only for every file system in the pool will be encrypted.
- H. You cannot create an encrypted file system in a zpool that was created prior to Oracle Solaris11. Create a new zpool in Solaris11, create an encrypted ZFS file system in the new zpool, and move or copy the data from the existing file system into the new encrypted file system.

Answer: ABCF

Explanation:

A (not H): You can use your existing storage pools as long as they are upgraded. You have the flexibility of encrypting specific file systems.

B (not E): Can I enable encryption on an existing pool?

Yes, the pool must be upgraded to pool version 30 to allow encrypted ZFS file systems and volumes.

C (not D): ZFS encryption is integrated with the ZFS command set. Like other ZFS operations, encryption operations such as key changes and rekey are performed online. F (not G): Encryption is the process in which data is encoded for privacy and a key is needed by the data owner to access the encoded data. You can set an encryption policy when a ZFS dataset is created, but the policy cannot be changed.

NEW QUESTION 216

User1 is attempting to assist user2 with terminating user2's process 1234. User1 entered the following: kill -9 1234 Why does the process continue to run?

- A. You can kill a process only if you are root.
- B. You can kill only a process that you own.
- C. You can kill the process only with the pkill command.
- D. You need to kill the process with a stronger kill signal.

Answer: B

Explanation:

Kill -9

Kill (terminates without cleanup)

Only works if issued by process owner or super user (root) The program cannot respond to this signal; it must terminate

Note: Unix provides security mechanisms to prevent unauthorized users from killing other processes. Essentially, for a process to send a signal to another, the owner of the signaling process must be the same as the owner of the receiving process or be the superuser.

NEW QUESTION 221

.....

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/>