



**Oracle**

**Exam Questions 1Z0-821**

Oracle Solaris 11 System Administrator

**NEW QUESTION 1**

Review the boot environments displayed on your system:

BE	Active	Mountpoint	Space	Policy	Created
oldBE	-	-	149.0K	static	2011-11-28 15:15
newBE	-	-	363.05M	static	2011-11-28 14:47
solaris	-	-	100.68M	static	2011-11-20 18:09
solaris-1	NR	/	19.07G	static	2012-01-22 07:23

Which option describes the solaris-1 BE?

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

**Answer: E**

**Explanation:**

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

Example:

Display your existing BE information.


# beadm list

BE Active Mountpoint Space Policy Created

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

**NEW QUESTION 2**

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



Identify the correct procedure for breaking the /prod\_data mirror, removing c4t0d0 and c4t2d0, and making the data on c4t0d0 and c4t2d0 accessible under the dev\_data mount point.

- A. zpool split pool1 pool2 c4t0d0 c4t2d0 zpool import pool2 zfs set mountpoint = /dev\_data pool2/prod\_data
- B. zpool detach pool1 pool2 zpool attach pool2 zfs set mountpoint = /dev\_data pool2/prod\_data
- C. zpool split pool1/prod\_data -n pool2/dev\_data zfs set mountpoint = /dev\_data pool2/prod\_data
- D. zpool split pool1 pool2 c4t0d0 c4t2d0 zpool import pool2

**Answer: D**

**Explanation:**

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

After the split operation, import the new pool.

### NEW QUESTION 3

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 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 5

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

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

**Answer:** B

#### Explanation:

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

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

### NEW QUESTION 6

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 7**

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 8**

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

```
pgrep userprogram l5317
```

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

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

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

**Answer: B**

**Explanation:**

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

To terminate a process:

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

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

**NEW QUESTION 9**

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

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

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

**Answer: A**

**NEW QUESTION 10**

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

Choose the command used to connect to `testzone`'s console.

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

**Answer: E**

**NEW QUESTION 10**

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

```
export dennis_says
```

You have also edited `/etc/skel/local.profile` to include the line: `dennis_says=world`  
You now create a new user account brian, and specify use of the bash shell. When brian logs in and enters  
`Echo $dennis_says`  
What will he see, and why?

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

**Answer:** A

**Explanation:**

The `$HOME/.profile` file is an initialization file that is executed after the `/etc/profile` when logging in to the Bourne or Korn shell. The file contains user preferences for variable settings. If the `ENV` variable is set to `.kshrc`, the `.kshrc` file executes every time a new shell begins execution. The `$HOME/.profile` is copied from the `/etc/skel/local.profile` file by the Administration Tool when creating a new account.  
Note: `/etc/skel/local.profile`  
Per-system configuration file for `sh/ksh/ksh93/bash` login sessions, installed for new users

**NEW QUESTION 15**

\_\_\_\_\_ 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 18**


A datalink can best be described as .

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

**Answer:** D

**Explanation:**

The command `dladm` is used to configure data-link interfaces in Sun Solaris. A configured data-link is represented in the system as interface that can be used for TCP/IP. Each data-link relies on either a single network device or an link aggregation device to send & receive packets.  
Network interfaces provide the connection between the system and the network. These interfaces are configured over data links, which in turn correspond to instances of hardware devices in the system.  
In the current model of the network stack, interfaces and links on the software layer build on the devices in the hardware layer. More specifically, a hardware device instance in the hardware layer has a corresponding link on the data-link layer and a configured interface on the interface layer. This one-to-one relationship among the network device, its data link, and the IP interface is illustrated in the figure that follows.  
Network Stack Showing Network Devices, Links, and Interfaces:



**NEW QUESTION 23**

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

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

**Answer:** DE

**Explanation:**

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

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

Examples of NVRAM variables:

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

diag-device net The diagnostic startup source device.

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

**NEW QUESTION 26**

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

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

Response:

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

Which command should you enter next?

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

**Answer:** B

**Explanation:**

Enable-if -t interface

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

-t, --temporary

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

**NEW QUESTION 29**

Review the information taken from your server:

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

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

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

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

**Answer:** A

**Explanation:**

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

-r

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

**NEW QUESTION 34**

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

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

**Answer:** B

**Explanation:**

Comparing sysidcfg File Keywords to System Configuration Profile Directives

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

sysidcfg File Keyword

System Configuration Profile Directives Etc.

**NEW QUESTION 36**

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 41**

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

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

**Answer:** C

**Explanation:**

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

-c class

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

-g

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

**NEW QUESTION 42**

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 47**

Which two capabilities are provided by the OpenBoot PROM?

- A. a command to safely shut down the system
- B. hardware testing and initialization
- C. booting from a disk or network
- D. starting the GRUB loader

**Answer:** BC

**Explanation:**

OpenBoot firmware is executed immediately after you turn on your system. The primary tasks of OpenBoot firmware are to:

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

**NEW QUESTION 51**

Your system is assigned an IP address object 192.168.0.222/24. However, the net mask — expressed as four octets — is required. Which is the correct netmask?

- A. 255.0.0.0
- B. 255.255.0.0
- C. 255.255.255.0
- D. 255.255.255.24
- E. 255.255.255.255

**Answer:** C

**Explanation:**

A 24-bit network mask is expressed as 255.255.255.0.

**NEW QUESTION 55**

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

```

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

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

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

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

```

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

**Answer:** B

**Explanation:**

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

For example:

```

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

```

**NEW QUESTION 60**

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  
 K3eropfqsMPPY6uiylh9vpiFX2r1LHcbx139+vG6HOtVvuS8+lfMDtawm3WQvRuOopz3vVy  
 5GtMwtaOgehsXoT930Ryev1bH5myPtWKlipITsOd2sX9k3tvjrmme4KCGGss=

#### NEW QUESTION 64

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. A virtual router
- C. A virtual switch
- D. A virtual bridge.
- E. A virtual network interface
- F. Nothing because a virtual switch is automatically created then the virtual network interfaces are created.

**Answer:** A

#### Explanation:

Etherstubs are pseudo ethernet NICs which are managed by the system administrator. You can create VNICs over etherstubs instead of over physical links. VNICs over an etherstub become independent of the physical NICs in the system. With etherstubs, you can construct a private virtual network that is isolated both from the other virtual networks in the system and from the external network. For example, you want to create a network environment whose access is limited only to your company developers than to the network at large. Etherstubs can be used to create such an environment.

Note: Oracle Solaris 11 introduces a new and powerful network stack architecture which includes:

- \* Networking virtualization with virtual network interface cards (VNICs) and virtual switching (etherstubs)
- \* Tight integration with zones
- \* Network resource management - efficient and easy to manage integrated quality of service (QoS) to enforce bandwidth limit on VNICs and traffic flows

#### NEW QUESTION 65

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 68

Given:

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

Which two commands will be successful?

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

**Answer:** CF

#### Explanation:

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

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

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

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

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

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

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

#### NEW QUESTION 71

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

- A. There are boot environments for zones.
- B. Administrators can delegate common administration tasks by using RBAC.
- C. Oracle Solaris 11 supports Solaris 8, 9, and 10 branded zones.
- D. You can migrate a physical Solaris 10 system and its non-global zones to a solaris10 branded zone running on an Oracle Solaris 11 system.
- E. It is possible to change the host ID of a zone.

**Answer:** ABD

#### Explanation:

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

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

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

#### NEW QUESTION 76

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

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


**Answer:** AC

#### Explanation:

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

#### NEW QUESTION 78

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 with `dladm create-vnic -z zone3 vnic1`.
- C. The virtual interface must be specified with `dladm create-vnic -z zone3 vnic1`.
- D. The virtual interface must be created with `ipadm create-vnic -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 82

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

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

**Answer:** AB

#### Explanation:

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

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

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

B: iSCSI targets cannot be configured as dump devices.

### NEW QUESTION 87

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/R00T/solaris	-	-	26.16M	static	2011-11-28 11:09
rpool/R00T/solaris/var	-	-	69.0K	static	2011-11-28 13:49
rpool/R00T/solaris/var@2011-11-28-18:49:38	-	-	0	static	2011-11-28 14:09
rpool/R00T/solaris/var@2011-11-28-19:09:23	-	-	975.0K	static	2011-11-28 12:29
rpool/R00T/solaris/var@install	-	-	70.0K	static	2011-11-28 13:49
rpool/R00T/solaris@2011-11-28-18:49:38	-	-	0	static	2011-11-28 14:09
rpool/R00T/solaris@2011-11-28-19:09:23	-	-	929.5K	static	2011-11-28 12:29
rpool/R00T/solaris@install	!R	-	2.0K	static	2011-11-28 13:49
solaris-1	-	-	1.0K	static	2011-11-28 13:49
rpool/R00T/solaris-1	-	-	-	-	-
rpool/R00T/solaris-1/var	-	-	57.0K	static	2011-11-28 14:09
z1BE	-	-	1.0K	static	2011-11-28 14:09
rpool/R00T/z1BE	-	-	-	-	-
rpool/R00T/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](http://docs.oracle.com/cd/E23824_01/html/E21801/unbootable.html#scrolltoc)

### NEW QUESTION 89

The /usr/bin/p7zip file that is part of the p7zip package has been overwritten. This server is critical to production and cannot be rebooted. Identify the command that would restore the file without requiring a reboot.

- A. pkg verify p7zip
- B. pkg fix p7zip
- C. pkg rebuild-index p7zip
- D. pkg revert p7zip
- E. pkg uninstdll p7zip
- F. pkg install p7zip
- G. pkg install --no-backup-be p7zip
- H. pkg refresh p7zip

**Answer:** D

**Explanation:**

Use the pkg revert command to restore files to their as-delivered condition.

**NEW QUESTION 91**

You have been asked to terminate a process that appears to be hung and will not terminate. The process table is shown below:  
 root 15163 15156 0 12:51:15 pts/3 0:00 hungscript What command will terminate the process?

- A. kill -9 15163
- B. kill -1 15163
- C. kill -15 15163
- D. kill -2 15163

**Answer:** A

**Explanation:**

Here we should use SIGTERM to terminate the process. Note:

When no signal is included in the kill command-line syntax, the default signal that is used is

-15 (SIGKILL). Using the -9 signal (SIGTERM) with the kill command ensures that the process terminates promptly. However, the -9 signal should not be used to kill certain processes, such as a database process, or an LDAP server process. The result is that data might be lost.

Tip - When using the kill command to stop a process, first try using the command by itself, without including a signal option. Wait a few minutes to see if the process terminates before

using the kill command with the -9 signal.

**NEW QUESTION 95**

Which two accurately describe the Solaris IPS repository?

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

**Answer:** BE

**Explanation:**

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

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

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

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

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

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

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

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

**NEW QUESTION 99**

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

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

Which three statements are correct?

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

**Answer:** DEF

**NEW QUESTION 103**

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

```
PUBLISHERTYPESTATUSURI
```

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

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

- A. On the server: set the pkg/inst\_root and pkg/readonly properties for the svc:/application/pkg/server:default service and enabled the service

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

**Answer:** AE

**Explanation:**

A: Configure the Repository Server Service

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

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

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

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

Restart the pkg.depotd repository service.

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

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

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

**NEW QUESTION 106**

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 107**

The COMSTAR framework provides support for the iSCSI protocol. Select three options that correctly describe the COMSTAR framework.

- A. iSCSI devices can be used as dump devices.
- B. SCSI commands are carried over IP networks and enable you to mount disk devices from across the network onto your local system.
- C. Large amounts of data can be transferred over an IP network with very little network degradation.
- D. COMSTAR allows you to convert any Solaris11 host into a SCSI target device that can be accessed over a storage network.
- E. One IP port can handle multiple iSCSI target devices.

**Answer:** BDE

**Explanation:**

B: By carrying SCSI commands over IP networks, the iSCSI protocol enables you to access block devices from across the network as if they were connected to the local system. COMSTAR provides an easier way to manage these iSCSI target devices.

D: Common Multiprotocol SCSI TARget, or COMSTAR, a software framework that enables you to convert any Oracle Solaris 11 host into a SCSI target device that can be accessed over a storage network by initiator hosts.

E: One IP port can handle multiple iSCSI target devices.

**NEW QUESTION 108**

Which modification needs to be made to the Service Management Facility before you publish a new package to the IPS repository?

- A. The pkg.depotd must be disabled.
- B. The pkg/readonly property for the application/pkg/server service must be set to false.
- C. The Pkg/writable\_root property for the application/Pkg/server service must be set to true.
- D. The pkg/image.root property for the application/pkg/server service must be set to the location of the repository.

**Answer:** D

**Explanation:**

pkg/image\_root

(astring) The path to the image whose file information will be used as a cache for file data.

**NEW QUESTION 113**

user1 has a disk quota of 0.5 MB. The user attempts to run the following command on a file called .bigfile that is 495 KB in size:

```
cp bigfile /tmp
```

Will the command execute successfully?

- A. Ye
- B. Quotas do not include any of the system files such as /tmp /swap.
- C. Ye

- D. The quota is set at the directory level, not the user level.
- E. N
- F. The command will fail because it will cause him to exceed his user quota.
- G. N
- H. A user cannot place files into the /tmp directory.

**Answer:** A

**Explanation:**

UFS quotas enable system administrators to control the size of file systems. Quotas limit the amount of disk space and the number of inodes, which roughly corresponds to the number of files, that individual users can acquire. For this reason, quotas are especially useful on the file systems where user home directories reside. As a rule, the public and /tmp file systems usually do not benefit significantly by establishing quotas. Note: The cp command copies files and directories.

**NEW QUESTION 117**

View the Exhibit.

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

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

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

**Answer:** D

**Explanation:**

ai\_manifest  
 - Automated installation manifest file format  
 Synopsis  
 /usr/share/install/ai.dtd.1  
 Some customizations have been made, such as the selection of specific locales.

### NEW QUESTION 119

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

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

**Answer:** ABDFH

#### Explanation:

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

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

```
$ sudo pkg install slim_install
```

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

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

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

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

### NEW QUESTION 120

United States of America export laws include restrictions on cryptography.

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

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

**Answer:** BC

#### Explanation:

B: Binary Signatures for Third-Party Software

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

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

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

### NEW QUESTION 121

You display the IP interface information with `ipmpstat -i`.

Which two characteristics are indicated by characters that may be included in the FLAGS column?

- A. default route
- B. IP forwarding enabled IS
- C. allocated to global zone
- D. unusable due to being inactive
- E. nominated to send/receive IPv4 multicast for its IPMP group

**Answer:** DE

#### Explanation:

FLAGS

Indicates the status of each underlying interface, which can be one or any combination of the following:

(D) d indicates that the interface is down and therefore unusable.

(E) M indicates that the interface is designated by the system to send and receive IPv6 multicast traffic for the IPMP group.

Note:

i indicates that the INACTIVE flag is set for the interface. Therefore, the interface is not used to send or receive data traffic.

s indicates that the interface is configured to be a standby interface.

m indicates that the interface is designated by the system to send and receive IPv4 multicast traffic for the IPMP group.

b indicates that the interface is designated by the system to receive broadcast traffic for the IPMP group.

h indicates that the interface shares a duplicate physical hardware address with another interface and has been taken offline. The h flag indicates that the interface is unusable.

### NEW QUESTION 125

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

Which four tasks are directly controlled by the OpenBoot firmware?

- A. Provides a list of boot entries from which to choose
- B. Allows hardware to identify itself and load its own plug-in device driver

- C. Loads the boot loader from the configured boot device
- D. Performs basic hardware testing
- E. Installs the console
- F. Reads and executes the boot archive
- G. Extract and executes the kernel image

**Answer:** ABCD

**Explanation:**

OpenBoot firmware is executed immediately after you turn on your system. The primary tasks of OpenBoot firmware are to:

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

Some notable features of OpenBoot firmware.

- \* Plug-in Device Drivers (B)

Most common tasks that you perform using OpenBoot

- \* Booting Your System

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

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

**NEW QUESTION 126**

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

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

Which four correctly describe the output?

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

**Answer:** ACEG

**Explanation:**

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

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

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

Note: Identifying ZFS Snapshot Differences (zfs diff)

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

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

File or Directory Change Identifier

- \* File or directory is modified or file or directory link changed M
- \* File or directory is present in the older snapshot but not in the newer snapshot

—

- \* File or directory is present in the newer snapshot but not in the older snapshot

+

- \* File or directory is renamed R

**NEW QUESTION 131**

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

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

**Answer:** AF

**Explanation:**

The Internet Control Message Protocol (ICMP) is one of the core protocols of the Internet Protocol Suite.

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

**NEW QUESTION 132**

Review the storage pool information:

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

Choose the correct procedure to repair this storage pool.

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

**Answer: C**

**Explanation:**

You might need to replace a disk in the root pool for the following reasons: The root pool is too small and you want to replace it with a larger disk. The root pool disk is failing. In a non-redundant pool, if the disk is failing so that the system won't boot, you'll need to boot from an alternate media, such as a CD or the network, before you replace the root pool disk. In a mirrored root pool configuration, you might be able to attempt a disk replacement without having to boot from alternate media. You can replace a failed disk by using the zpool replace command. Some hardware requires that you offline and unconfigure a disk before attempting the zpool replace operation to replace a failed disk.

For example:

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

**NEW QUESTION 133**

The following line is from /etc/shadow in a default Solaris 11 Installation:  
jack: \$5\$9JFrt54\$7JdwmO.F11Zt/jFeeOhDmnw93LG7Gwd3Nd/cwCcNWFFg:0:15:30:3:: Which two are true?

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

**Answer: BE**

**Explanation:**

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

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

**NEW QUESTION 136**

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

- A. all service transition states
- B. service dependencies that have stopped or faulted
- C. service configuration modifications
- D. legacy services that have not started
- E. services that have been disabled
- F. service fault management events
- G. processes that have been killed

**Answer: AEFB**

**Explanation:**

Note 1: State Transition Sets are defined as: to<state>

Set of all transitions that have <state> as the final state of the transition.

form-<state>

Set of all transitions that have <state> as the initial state of the transition.

<state>

Set of all transitions that have <state> as the initial state of the transitional. Set of all transitions. (A)

Valid values of state are maintenance, offline (G), disabled (E), online and degraded. An example of a transitions set definition: maintenance, from-online, to-degraded.

F: In this context, events is a comma separated list of SMF state transition sets or a comma separated list of FMA (Fault Management Architecture) event classes. events cannot have a mix of SMF state transition sets and FMA event classes. For convenience, the tags problem- {diagnosed, updated, repaired, resolved} describe the lifecycle of a problem diagnosed by the FMA subsystem - from initial diagnosis to interim updates and finally problem closure.

Note 2:

SMF allows notification by using SNMP or SMTP of state transitions. It publishes Information Events for state transitions which are consumed by notification daemons like snmp-notify(1M) and smtp-notify(1M). SMF state transitions of disabled services do not generate notifications unless the final state for the transition is disabled and there exist notification parameters for that transition. Notification is not be generated for transitions that have the same initial and final state.

**NEW QUESTION 140**

User jack logs in to host solaris and then attempts to log in to host oracle using ssh. He receives the following error message:  
 The authenticity of host oracle (192.168.1.22) can't be established. RSA key fingerprint is 3B:23:a5:6d:ad:a5:76:83:9c:c3:c4:55:a5:18:98:a6  
 Are you sure you want to continue connecting (yes/no)?  
 Which two are true?

- A. The public host key supplied by solaris is not known to the host oracle.
- B. The error would not occur if the RSA key fingerprint shown in the error message was added to the /etc/ssh/known\_hosts file on solaris.
- C. The private host key supplied by oracle is not known to solaris.
- D. If jack answers yes, the RSA public key for the host oracle will be added to the known\_hosts file for the user jack.
- E. The public host key supplied by oracle is not known to the host solaris.

**Answer:** BD

**Explanation:**

The fingerprints are used to guard against man in the middle attacks. Since ssh logins usually work over the internet (an insecure connection), someone could hijack your connection. When you try to log into yourmachine.com, he could get "in the middle" and return your challenge as if he was yourmachine.com. That way, he could get hold of your login password.

To make this attack harder, ssh stores the fingerprint of the server's public key on the first connection attempt. You will see a prompt like:  
 The authenticity of host 'eisen (137.43.366.64)' can't be established. RSA key fingerprint is cf:55:30:31:7f:f0:c4:a0:9a:02:1d:1c:41:cf:63:cf. Are you sure you want to continue connecting (yes/no)

When you enter yes, ssh will add the fingerprint to your known\_hosts file. you will see

Code:  
 Warning: Permanently added 'eisen, 137.43.366.64' (RSA) to the list of known hosts.

The next time you login, ssh will check whether the host key has changed. A changing host key usually indicates a man in the middle attack, and ssh refuses to connect.

**NEW QUESTION 142**

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 146**

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

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

NAME	STATE	READ	WRITE	CKSUM
pool1	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
c3t3d0	ONLINE	0	0	0
c3t4d0	ONLINE	0	0	0
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 c3t4do.
- B. The data is striped across disks c3t3d0 and c3t4do 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 148

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 150

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 151

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 155**

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 157**

.....

## Thank You for Trying Our Product

### We offer two products:

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

2nd - Questions and Answers in PDF Format

### 1Z0-821 Practice Exam Features:

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

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