



Oracle

Exam 1z0-821

Oracle Solaris 11 System Administration

Version: 12.0

[Total Questions: 234]

Topic break down

Topic	No. of Questions
Topic 1: Volume A	100
Topic 2: Volume B	134

Topic 1, Volume A

Question No : 1 - (Topic 1)

Which two options are accurate regarding the non-global zone console?

- A. Access the non-global zone console by using the `zlogin -c` command.
- B. Access the non-global zone console by using the `zlogin -1` command.
- C. Disconnect from the non-global zone console by using the `~.` keys.
- D. Disconnect from the non-global zone console by using the `#.` keys.

Answer: A,C

Explanation:

A: How to Log In to the Zone Console Use the `zlogin` command with the `-C` option and the name of the zone,for example,my-zone.

```
global# zlogin -C my-zone
```

C: To disconnect from a non-global zone,use one of the following methods.

* To exit the zone non-virtual console:

```
zonename# exit
```

* To disconnect from a zone virtual console,use the tilde (`~`) character and a period:

```
zonename# ~.
```

Question No : 2 - (Topic 1)

View the following information for a software package:

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

Which command would you use to display this information for a software package that is

not currently installed on your system?

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

Answer: B

Explanation:

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

Example:

Listing Information About an Uninstalled Package

```
# pkg info -r php-52
```

Name: web/php-52

Summary: PHP Server 5.2

Description: PHP Server 5.2

Category: Development/PHP

State: Not Installed

Publisher: solaris

Version: 5.2.17

Build Release: 5.11

Branch: 0.175.0.0.0.1.530

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

Size: 44.47 MB

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

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

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

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

Question No : 3 - (Topic 1)

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
```

Question No : 4 - (Topic 1)

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

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

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

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

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

- ✍ The server file systems must remain available at all times and can never be taken offline.
- ✍ The backup must not only provide for the recovery of the most recent version of a file, but must also allow recovery of previous versions of a file created since Saturday's full backup.

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

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

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

Question No : 5 - (Topic 1)

Your SPARC server will not boot into multi user-server milestones and you need to troubleshoot to out why. You need to start the server with minimal services running so that you can go through each milestone manually to troubleshoot the issue.

Select the option that boots the server with the fewest services running.

- A. boot -s
- B. boot milestone none
- C. boot -m milestone=single-user
- D. boot -m milestone=none
- E. boot -m none

Answer: D

Explanation:

The command boot -m milestone=none is useful in repairing a system that have problems booting early.

Boot Troubleshooting:

To step through the SMF portion of the boot process,start with:

```
boot -m milestone=none
```

Then step through the milestones for the different boot levels:

```
svcadm milestone svc:/milestone/single-user:default
```

```
svcadm milestone svc:/milestone/multi-user:default
```

```
svcadm milestone svc:/milestone/multi-user-server:default
```

Question No : 6 - (Topic 1)

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

```
1:07am up 346 day(s),12:03,4 users,load average: 0.02,0.02,0.01
```

Which command will be useful in your next step of troubleshooting?

- A. ipadm
- B. traceroute
- C. dladm

- D. snoop
- E. arp

Answer: B

Explanation:

Test the remote connection with traceroute.

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

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

Question No : 7 - (Topic 1)

Examine the following command sequence:

```
jack@solaris:~$ cd
jack@solaris:~$ ls -ld .
drwxr-xr-- 1 jack other 23 2012-01-04 22:12 .
jack@solaris:~$ ls -ld dira
drwxr-xr-x 3 jack other 3 2012-01-06 08:27 dira
jack@solaris:~$ ls -lR dira
dira:
total 2
drwxrwxr-t 2 jack other 3 2012-01-06 08:27 dirb

dira/dirb:
total 1
-rw-r--rw- 1 jack other 8768 2012-01-06 08:28 testfile
-rwxr-xr-x 1 jack other 3398 2011-12-30 12:10 scriptfile
```

Which three are true?

- A. User jill, a member of the group other, can do a long listing (ls -l) of user jack's home

directory.

- B. User oracle, not a member of the group other, cannot execute the shell script scriptfile.
- C. User jill, a member of the group other, can make a copy of testfile in dira.
- D. User jill, a member of the group other, can edit the data content of testfile.
- E. User jill, a member of the group other, cannot delete testfile.

Question No : 8 - (Topic 1)

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

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

Which three statements are correct?

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

Answer: A,C,D

Explanation:

A: dir followed by a single letter.

C: dir followed by any characters ending with a.

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

only dirabc matches

Question No : 9 - (Topic 1)

You are the administrator of a system that a large number of developers work on. These developers crash the system,and their applications,on a regular basis.

What command would you use to configure where the core files are saved?

- A. savecore
- B. dumpadm
- C. svcadm
- D. proc
- E. coreadm

Answer: E

Explanation:

The coreadm command is used to specify the name and location of core files produced by abnormally-terminating processes.

Question No : 10 - (Topic 1)

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.

Question No : 11 - (Topic 1)

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 G8 total size (rounded to the nearest GB).
- E. The storage pool is 292 GB total size (rounded to the nearest GB).

Answer: A,E

Question No : 12 - (Topic 1)

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: B,E

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 that 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.

Question No : 13 - (Topic 1)

User jack, whose account is configured to use the Korn shell, logs in and examines the value of his PATH environment variable:

```
jack@solaris: echo $PATH
/usr/gnu/bin:/usr/bin:/usr/sbin:/sbin
```

There is a shell script in jack's home directory called useradd:

```
-r-xr-xr-x 2 jack other 1239 2012-01-05 11:42 useradd
```

While in his home directory, jack attempts to run the script:

```
jack@solaris: useradd
```

What will happen, and why?

- A. He will get a "file not found" error, because the current directory is not in his search path.
- B. He will get a "file not found" error, because his home directory is not in his search path.
- C. The useradd script will execute, because jack is in the same directory that the script is located in.
- D. The command /user/sbin/useradd will execute, because it is the last match in the search path.
- E. The command /user/sbin/useradd will execute, because it is the first match in the search path.

Answer: D

Question No : 14 - (Topic 1)

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).

Question No : 15 - (Topic 1)

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

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

D. zonecfg –icv dbzone info

Answer: C

Explanation:

zonecfg info

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

Note:

zonecfg –z

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

Incorrect answer:

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

list option:

Display the name of the current zones, or the specified zone if indicated.

B: No such command.

D: no such options zonecfg –icv

Question No : 16 - (Topic 1)

View the Exhibit and review the file system information displayed from a remote server.

```

Question Exhibit
root@solaris:~# df -h
Filesystem      Size  Used Available Capacity  Mounted on
rpool/ROOT/solaris 16G  3.4G  9.8G    26%    /
/devices        OK    OK    OK      0%    /devices
/dev            OK    OK    OK      0%    /dev
ctfs            OK    OK    OK      0%    /system/contract
proc           OK    OK    OK      0%    /proc
mnttab         OK    OK    OK      0%    /etc/mnttab
swap          1.1G  1.4M  1.1G    1%    /system/volatile
objfs         OK    OK    OK      0%    /system/object
sharefs       OK    OK    OK      0%    /etc/dfs/sharetab
/usr/lib/libc/libc_hwcap1.so.1
fd            13G  3.4G  9.8G    26%    /lib/libc.so.1
rpool/ROOT/solaris/var
swap          16G  183M  9.8G    2%    /var
rpool/export   1.1G  48K   1.1G    1%    /tmp
rpool/export/home 16G  32K   9.8G    1%    /export/home
rpool/export/home/bcalkins
rpool         16G  714K  9.8G    1%    /export/home/bcalkins
pool1/data    16G  39K   9.8G    1%    /rpool
root@solaris:~#

```

You are configuring a new server. This new server has the following storage pool configured:

NAME	SIZE	ALLOC	FREE	CAP	DEDUP	HEALTH	ALTROOT
Pool1	15.9G	85K	15.9G	0%	1.00x	ONLINE	-

This new server also has the following file systems configured:

NAME	USED	AVAIL	REFER	MOUNTPOINT
pool1	85K	15.6G	31K	/pool1
rpool	5.81G	9.82G	39K	/rpool
rpool/ROOT	3.82G	9.82G	31K	legacy
rpool/ROOT/solaris	3.82G	9.82G	3.40G	/
rpool/ROOT/solaris/var	333M	9.82G	183M	/var
rpool/dump	970M	9.85G	940M	-
rpool/export	796K	9.82G	32K	/export
rpool/export/home	764K	9.82G	32K	/export/home
rpool/export/home/bcalkins	714K	9.82G	714K	/export/home/bcalkins
rpool/swap	1.03G	9.85G	1.00G	-

When you are finished building this new server, the pool1/data dataset must be an exact duplicate of note server. What is the correct procedure to create the pool1/data dataset on this new server?

- A. zfs create -o mountpoint=/data -o refquota=1g pool1/data
- B. zfs set mountpoint=none pool1 zfs create pool1/data
- C. zfs set mountpoint=none pool1 zfs create -o mountpoint=/data -o quota=1g pool1/data
- D. zfs create quota=1g pool1/data
- E. zfs create mountpoint=/data pool1/data
- F. zfs set quota=1g pool1/data

Answer: A

Question No : 17 - (Topic 1)

You are attempting to edit your crontab file in the bash shell. Instead of getting your usual vi interface, you are presented with an unfamiliar interface. In order to have your editor of choice-vi- what command must you type after exiting the unfamiliar editor?

- A. EDITOR=vi
- B. crontab=vi
- C. crontab - e vi
- D. env

Answer: A

Explanation:

Set the EDITOR variable to vi.

Commands like `crontab -e` will use ed per default. If you'd like to use some better editor (like vi) you can use the environment variable EDITOR:

```
# EDITOR=vi; crontab -e will open the users crontab in vi. Of course you can set this variable permanently.
```

Incorrect answers

C: -e Edits a copy of the current user's crontab file, or creates an empty file to edit if crontab does not exist. When editing is complete, the file is installed as the user's crontab file. If a username is given, the specified user's crontab file is edited, rather than the current user's crontab file; this can only be done by a user with the solaris.jobs.admin authorization. The environment variable EDITOR determines which editor is invoked with the -e option. The default editor is ed(1). All crontab jobs should be submitted using crontab. Do not add jobs by just editing the crontab file, because cron is not aware of changes made this way.

Question No : 18 - (Topic 1)

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 platform. The feature is supported only on the SPARC platform.
- B. Use the iscsitadm command on the x86 platform when configuring an iSCSI target.
- C. Install the storage-server group package on this system.
- D. 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.

Question No : 19 - (Topic 1)

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.

Question No : 20 - (Topic 1)

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

- A)

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

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

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

```
          capacity      operations      bandwidth
pool      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
```

Question No : 21 - (Topic 1)

Choose three options that describe the features associated with a Live Media installation.

- A. does not allow the root user to log in to the system directly from the console (or any terminal)
- B. provides a "hands free" installation
- C. installs the desktop based packages
- D. can be used to install only x86 platforms
- E. installs the server-based set of packages only
- F. allows both automatic and manual configuration of the network

G. installs both the server-based and desktop-based package

Answer: B,C,D

Explanation:

The graphical installer is officially known as the "Live Media." This means that Oracle Solaris can be booted into RAM, causing zero impact on your existing operating system. After it is loaded, you are free to experiment with Oracle Solaris to determine whether it is something you would like to install to your system.

You can download Oracle Solaris 11 Live Media for x86, which is an approximately 800 MB image file, and use a DVD burner to create the disk, or you can use the ISO image directly in a virtual machine or through the Oracle Integrated Lights Out Manager (ILOM) Remote Console.

The Live Media is not intended for long-term use. For example, any changes that you make to the system are lost when the system is shut down. Therefore, the next logical step is to install Oracle Solaris on the system, which the Live Media makes easy by placing an Install Oracle Solaris icon right on the desktop. But before we head down that road, let's step back a bit and consider the installation options.

Note: The Live Media provides administrators with an opportunity to explore the Oracle Solaris 11 environment without installing it on a system. The system boots off the media directly allowing administrators to start the installer should they choose to install it to a system.

Question No : 22 - (Topic 1)

Identify the correct description of an IPS image.

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

Answer: D

Explanation:

An image is a location where packages can be installed.

An image can be one of three types:

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

Question No : 23 - (Topic 1)

You have installed an update to the gzip package and need to "undo" the update and return the package to its "as-delivered" condition. Which command would you use?

- A. pkg undo
- B. pkg revert
- C. pkg fix
- D. pkg uninstall

Answer: B

Explanation:

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

Question No : 24 - (Topic 1)

How are operating system updates distributed in the Oracle Solaris 11 environment?

- A. Updates are only available to customers with an active support contract. The updates are distributed through the My Oracle Support web portal and installed in a central location. All software packages are then updated manually from the command line using the smpatch command.
- B. Patches are downloaded from <http://support.oracle.com> either automatically or manually. All software packages are then updated manually from the command line using the smpatch or patchadd commands.

C. Software updates are published as packages to a repository. All software packages are then updated manually from the command line using the pkg command.

D. Software updates, published as packages to an OS image. All software packages are then updated manually from the command line using the pkg command.

Answer: C

Explanation:

* Updating all of the packages on your installed system – To update all of the packages on your system that have available updates, use the pkg update command, as follows:

```
# pkg update
```

Running this command updates packages that you might not otherwise consider updating, for example, kernel components and other low-level system packages.

* Adding or updating individual packages – To add individual software packages, use the pkg install command. Any dependent packages are also updated at the same time.

* install package updates that deliver fixes – A pkg update operation might include bug fixes, so the operation is similar to applying a specific patch or patches in previous Oracle Solaris releases.

Note: The IPS interfaces first check for updates for currently installed packages before retrieving them via the network. By default, interfaces check repository catalogs in the following locations:

* The default installation repository at pkg.oracle.com/solaris/release.

* The support repository in My Oracle Support. This repository is restricted to users with Oracle Solaris 11 Express support contracts, and it contains packages with the latest bug fixes. For this reason, a support contract must be purchased for production deployments.

Question No : 25 - (Topic 1)

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.

Question No : 26 - (Topic 1)

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 BE. The install,update,uninstall,or revert operation is not performed

if a new BE is required.

C. If a BE is created,do not set it as the active BE on the next boot

D. A reboot is required for all package operations

E. A reboot is not required after a package operation.

F. For package operations that require a reboot,this policy creates a new BE set as active on the next boot.

Answer: D,F

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.

Question No : 27 - (Topic 1)

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

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

The server has two spare 146-GB disk drives: c3t5d0 c3t6d0

You need to add more space to the pool1 storage pool. Which command would add more mirrored storage to the pool1 storage pool?

- A. `zpool add pool1 mirror c3t5d0 c3t6d0`
- B. `zpool attach pool1 mirror c3t5d0 c3t6d0`
- C. `zpool attach pool1 c3r3d0 c3r5d0; zpool attach pool1 c3r4d0 c3r6d0`
- D. `zpool add pool1 c3r3d0 c3r5d0; zpool add pool1 c3r4d0 c3r6d0`

Answer: A

Question No : 28 - (Topic 1)

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.

Question No : 29 - (Topic 1)

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

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

Answer: A

Explanation:

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

cd -

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

Note:

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

Question No : 30 - (Topic 1)

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.

Question No : 31 - (Topic 1)

You log in to the system as user1,then switch user to root by using the su - command. After entering the correct password,yon enter the following commands:

whoami;who am i;id

Which option correctly represents the output?

- A) uid=0 (root) gid=0 (root)
user1 console Dec 30 20:20
root
- B) root
user1 console Dec 30 20:20
uid=0 (root) gid=0 (root)
- C) user1 console Dec 30 20:20
root
uid=0 (root) gid=0 (root)
- D) uid=0 (root) gid=0 (root)
root
user1 console Dec 30 20:20

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

Answer: B

Explanation:

* The whoami utility displays your effective user ID as a name.

Here this would be root.

* who am i

The command who shows who is logged on.

Here this would be:

user1 console Dec 30 20:20

* The id utility displays the user and group names and numeric IDs,of the calling process,to the standard output. If the real and effective IDs are different,both are displayed,otherwise only the real ID is displayed.

Here this would be:

uid=0(root) gid=0(root)

Note:

Each UNIX proces has 3 UIDs associated to it. Superuser privilege is UID=0.

Real UID

This is the UID of the user/process that created THIS process. It can be changed only if the running process has EUID=0.

Effective UID

This UID is used to evaluate privileges of the process to perform a particular action. EUID can be change either to RUID,or SUID if EUID!=0. If EUID=0,it can be changed to anything.

Saved UID

If the binary image file,that was launched has a Set-UID bit on,SUID will be the UID of the owner of the file. Otherwise,SUID will be the RUID.

Question No : 32 - (Topic 1)

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
```

Question No : 33 - (Topic 1)

You are asked to determine user jack's default login directory. Which command would provide you with useful information?

- A. `cat /etc/passwd | grep jack`
- B. `cat /etc/group | grep jack`
- C. `cat /etc/shadow | grep jack`
- D. `cat /etc/default/passwd | grep jack`

Answer: A

Explanation:

The `/etc/passwd` contains one entry per line for each user (or user account) of the system. All fields are separated by a colon (`:`) symbol. Total seven fields as follows.

1. Username: It is used when user logs in. It should be between 1 and 32 characters in length.
2. Password: An `x` character indicates that encrypted password is stored in `/etc/shadow` file.
3. User ID (UID): Each user must be assigned a user ID (UID). UID 0 (zero) is reserved for root and UIDs 1-99 are reserved for other predefined accounts. Further UID 100-999 are reserved by system for administrative and system accounts/groups.
4. Group ID (GID): The primary group ID (stored in `/etc/group` file)
5. User ID Info: The comment field. It allow you to add extra information about the users such as user's full name,phone number etc. This field use by `finger` command.
6. Home directory: The absolute path to the directory the user will be in when they log in. If this directory does not exists then users directory becomes `/`
7. Command/shell: The absolute path of a command or shell (`/bin/bash`). Typically,this is a shell. Please note that it does not have to be a shell.

Question No : 34 - (Topic 1)

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: B,E

Explanation:

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

* username: jack

* encrypted password: \$5\$9JFrt54\$7JdwmO.F11Zt/
jFeeOhDmnw93LG7Gwd3Nd/cwCcNFFFg

* 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

Question No : 35 - (Topic 1)

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: B,D

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.

Question No : 36 - (Topic 1)

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

What is the name of this utility?

- A. ppgsz

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

Answer: B

Explanation:

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

Question No : 37 - (Topic 1)

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

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

Answer: A,B

Explanation:

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

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

Question No : 38 - (Topic 1)

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 reservation=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
```

Question No : 39 - (Topic 1)

You created an IP address for interface `net.3` with the following command, which executed successfully:

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

You then ran:

```
ipadm show-if
```

The result indicated that the interface was down.

You then ran:

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

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

```
ipadm show-if
```

The last command indicated that the interface was up.

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

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

Answer: B

Explanation:

The first IP address is already in use.

Question No : 40 - (Topic 1)

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

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

201301(Pr201301)

What are the two problems with your plan?

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

Answer: C,D

Explanation:

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

D: Group names contain only lowercase characters and numbers.

Question No : 41 - (Topic 1)

User brian changes the permissions for db_data this command:

```
chmod 4755 db_data
```

What is true?

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

Answer: C

Explanation:

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

Here we do not know if brian owns db_data.

Note:

Permission

7 full

6 read and write

5 read and execute

4 read only

3 write and execute

2 write only
1 execute only
0 none

0 --- no permission
1 --x execute
2 -w- write
3 -wx write and execute
4 r-- read
5 r-x read and execute
6 rw- read and write
7 rwx read,write and execut

Solaris: Solaris Advanced User's Guide

Question No : 42 - (Topic 1)

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: D,E,F

Question No : 43 - (Topic 1)

Which three options accurately describe Oracle Solaris 11 zones?

- A. can be NFS servers
- B. are whole root type only
- C. cannot have their own time zone setting
- D. can execute zfs and zpool commands (from a non-global zone)
- E. are virtualized operating system environments,each with its own dedicated OS and kernel
- F. are virtualized operating system environments,created with a single instance of the OS shared kernel

Answer: A,D,F

Explanation:

A: Zones can use Oracle Solaris 11 products and features such as the following:

Oracle Solaris ZFS encryption

Network virtualization and QoS

CIFS and NFS

C: Non-global zones cannot modify the system clock by default,but each zone can have a separate time zone setting.

F (not E): The Oracle Solaris Zones partitioning technology is used to virtualize operating system services and provide an isolated and secure environment for running applications. A zone is a virtualized operating system environment created within a single instance of the Oracle Solaris operating system.

Question No : 44 - (Topic 1)

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: A,E

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.

Question No : 45 - (Topic 1)

You need to make sure that all of the software packages on your server are up to date. Without installing any updates, which two commands would display .my software updates that are available in the default Oracle repository?

- A. `pkg list -u`
- B. `pkg verify -u '*'`
- C. `pkg search -u`
- D. `pkg info -r '*'`
- E. `pkg install -nv`
- F. `pkg update -nv '*'`

Answer: A,D

Explanation:

A: the `pkg list` command displays a list of packages in the current image, including state and other information. By default, package variants for a different architecture or zone type are excluded.

D: `pkginfo` displays information about software packages that are installed on the system (with the first synopsis, with `-l`) or that reside on a particular device or directory (with the second synopsis, with `-r`).

Without options, `pkginfo` lists the primary category, package instance, and the names of all completely installed and partially installed packages. It displays one line for each package selected.

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

Question No : 46 - (Topic 1)

View the exhibit.

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

The configuration information in the exhibit is displayed on your system immediately after installing the OS.

Choose the option that describes the selection made during the Installation of the OS to obtain this configuration.

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

Answer: A

Explanation:

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

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

Question No : 47 - (Topic 1)

You need to know the IP address configured on interface net3, and that the interface is up. Which command confirms these?

- A. ipadm show-if
- B. ipadm up-addr
- C. ipadm show-addr
- D. ipadm enable-if
- E. ipadm refresh-addr
- F. ipadm show-addrprop

Answer: C

Explanation:

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

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

Example:

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

Question No : 48 - (Topic 1)

You need to install the solaris-desktop group package. Which command would you use to list the set of packages included in that software group?

- A. pkg search
- B. pkg info
- C. pkginfo
- D. pkg contents

Answer: A

Explanation:

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

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

Question No : 49 - (Topic 1)

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

```

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

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

```

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

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

A. Option A

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

Answer: C

Question No : 50 - (Topic 1)

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.

Question No : 51 - (Topic 1)

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

How can this be accomplished?

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

Answer: E

Explanation:

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

The value assigned by umask is subtracted from the default.

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

umask 000 would grant full permissions.

Note: 777 full permissions

Question No : 52 - (Topic 1)

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

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

Answer: B

Explanation:

The pkg info command provides detailed information about a particular IPS package.

Note: The pkginfo command does the same for any SVR4 packages you may have installed on the same system.

pkg info example:

```
$ pkg info p7zip
```

Name: compress/p7zip

Summary: The p7zip compression and archiving utility

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

Category: System/Core

State: Installed

Publisher: solaris

Version: 9.20.1

Build Release: 5.11

Branch: 0.175.0.0.0.2.537

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

Size: 6.73 MB

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

Question No : 53 - (Topic 1)

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 bootarguments from the default boot device

Note: boot has the following general format:

boot [device-specifier] [arguments]

where device-specifier and arguments are optional.

Question No : 54 - (Topic 1)

A user account must be a member of a primary group,and may also be a member of one or more secondary groups. What is the maximum total number of groups that one user can concurrently belong to?

- A. 15
- B. 16
- C. 17
- D. 63
- E. 64
- F. 65
- G. The number of groups one user can concurrently belong to is unlimited in Solaris 11.

Answer: B

Explanation:

Each user belongs to a group that is referred to as the user's primary group. The GID number,located in the user's account entry within the /etc/passwd file,specifies the user's primary group.

Each user can also belong to up to 15 additional groups,known as secondary groups. In the /etc/group file,you can add users to group entries,thus establishing the user's secondary group affiliations.

Note (4 PSARC/2009/542):

his project proposes changing the maximum value for NGROUPS_MAX from 32 to 1024 by changing the definition of NGROUPS_UMAX from 32 to 1024.

The use for a larger number of groups is described in CR 4088757, particular in the case of Samba servers and ADS clients; the Samba servers map every SID to a Unix group. Users with more than 32 groups SIDs are common. We've seen reports varying from "64 is enough", "128 is absolutely enough" and "we've users with more 190 group SIDS).

NGROUPS_MAX as defined by different Unix versions are as follows (<http://www.j3e.de/ngroups.html>):

Linux Kernel >= 2.6.3 65536

Linux Kernel < 2.6.3 32

Tru64 / OSF/1 32

IBM AIX 5.2 64

IBM AIX 5.3 ... 6.1 128

OpenBSD, NetBSD, FreeBSD, Darwin (Mac OS X) 16

Sun Solaris 7,8,9,10 16 (can vary from 0-32)

HP-UX 20

IRIX 16 (can vary from 0-32)

Plan 9 from Bell Labs 32

Minix 3 0 (Minix-vmd: 16)

QNX 6.4 8

Question No : 55 - (Topic 1)

What determines which bits in an IP address represent the subnet, and which represent the host?

- A. Subnet
- B. unicast
- C. netmask

- D. multicast
- E. broadcast

Answer: C

Explanation:

A subnetwork, or subnet, is a logically visible subdivision of an IP network. The practice of dividing a network into two or more networks is called subnetting.

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.

Question No : 56 - (Topic 1)

Consider the following commands:

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

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

Answer: B

Question No : 57 - (Topic 1)

View the Exhibit to see the information taken from the installation log file.

Based on the information presented in the Exhibit, which two options describe the state of the system when the server is booted for the first time after the installation is complete?

```
Disk: 16.0GB sata
Partition: 16.0GB Solaris2
Slice 0: 16.0GB rpool

Time Zone: US/Eastern

Language: *The following can be changed when logging in.
Default language: English

Keyboard layout: *The following can be changed when logging in.
Default keyboard layout: US-English

Terminal type: sun-color

Users:
No user account

Network:
Computer name: solaris
Network Configuration: None
```

- A. NWAM will be used to configure the network interface.
- B. The network/physical service is offline.
- C. You cannot log in from the console as root. You must first log in as a user and then su to root account.
- D. The root user can log in from the console login.
- E. You will be prompted to configure the network interface after the initial login.

Answer: B,D

Question No : 58 - (Topic 1)

Which three of the components could be used in a ZFS storage pool, but are not recommended configurations?

- A. A file on a UFS file system
- B. A Veritas Volume Manager (VxVM) volume
- C. A LUN In a hardware RAID array
- D. A disk slice from an SMI labeled disk
- E. A Solaris Volume Manager (SVM) volume
- F. An EFI labeled disk

Answer: A,B,E

Explanation:

A: ZFS also allows you to use UFS files as virtual devices in your storage pool. This feature is aimed primarily at testing and enabling simple experimentation, not for production use. The reason is that any use of files relies on the underlying file system for consistency. If you create a ZFS pool backed by files on a UFS file system, then you are implicitly relying on UFS to guarantee correctness and synchronous semantics.

However, files can be quite useful when you are first trying out ZFS or experimenting with more complicated layouts when not enough physical devices are present. All files must be specified as complete paths and must be at least 64 Mbytes in size.

B,E: You can construct logical devices for ZFS using volumes presented by software-based volume managers, such as Solaris Volume Manager (SVM) or Veritas Volume Manager (VxVM). However, these configurations are not recommended. While ZFS functions properly on such devices, less-than-optimal performance might be the result.

Question No : 59 - (Topic 1)

You are logged in as root to a newly installed Solaris 11 system. You issue the command `useradd -d`, and then examine the `/usr/sadm/defadduser` file. This file includes the entry `defshell=/bin/sh`. Which shell will now be the default for the next account created?

- A. bash shell
- B. C shell
- C. korn shod
- D. bourne shell

Answer: A

Explanation:

Oracle Solaris 11 introduces user environment and command-line argument changes that include the following:

- * Shell changes - The default shell, `/bin/sh`, is now linked to `ksh93`. The default user shell is the Bourne-again (bash) shell.
- * The legacy Bourne shell is available as `/usr/sunos/bin/sh`.
- * The legacy `ksh88` is available as `/usr/sunos/bin/ksh` from the `shell/ksh88` package.
- * Korn shell compatibility information is available in `/usr/share/doc/ksh/COMPATIBILITY`.

Question No : 60 - (Topic 1)

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.

Question No : 61 - (Topic 1)

_____ 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.

Question No : 62 - (Topic 1)

You have a user that needs to use the cron tool to schedule some repetitive tasks. When the user enters the crontab –e command in a terminal window,the following error appears:

```
crontab: you are not authorized to use cron. Sorry
```

In order to troubleshoot this issue,in what directory would you start your invest

- A. /etc/cron.d
- B. /var/spool/cron
- C. /var/spool/cron/crontable
- D. /var/spool/cron/atjobs

Answer: A

Explanation:

```
crontab: you are not authorized to use cron. Sorry.
```

This message means that either the user is not listed in the cron.allow file (if the file exists),or the user is listed in the cron.deny file.

You can control access to the crontab command by using two files in the /etc/cron.d directory: cron.deny and cron.allow. These files permit only specified users to perform crontab command tasks such as creating,editing,displaying,or removing their own crontab files.

The cron.deny and cron.allow files consist of a list of user names,one user name per line.

Question No : 63 - (Topic 1)

View the Exhibit.

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

Which is true regarding the disk drive?

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

Answer: A

Explanation:

Installing a ZFS Root Pool

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

Question No : 64 - (Topic 1)

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:

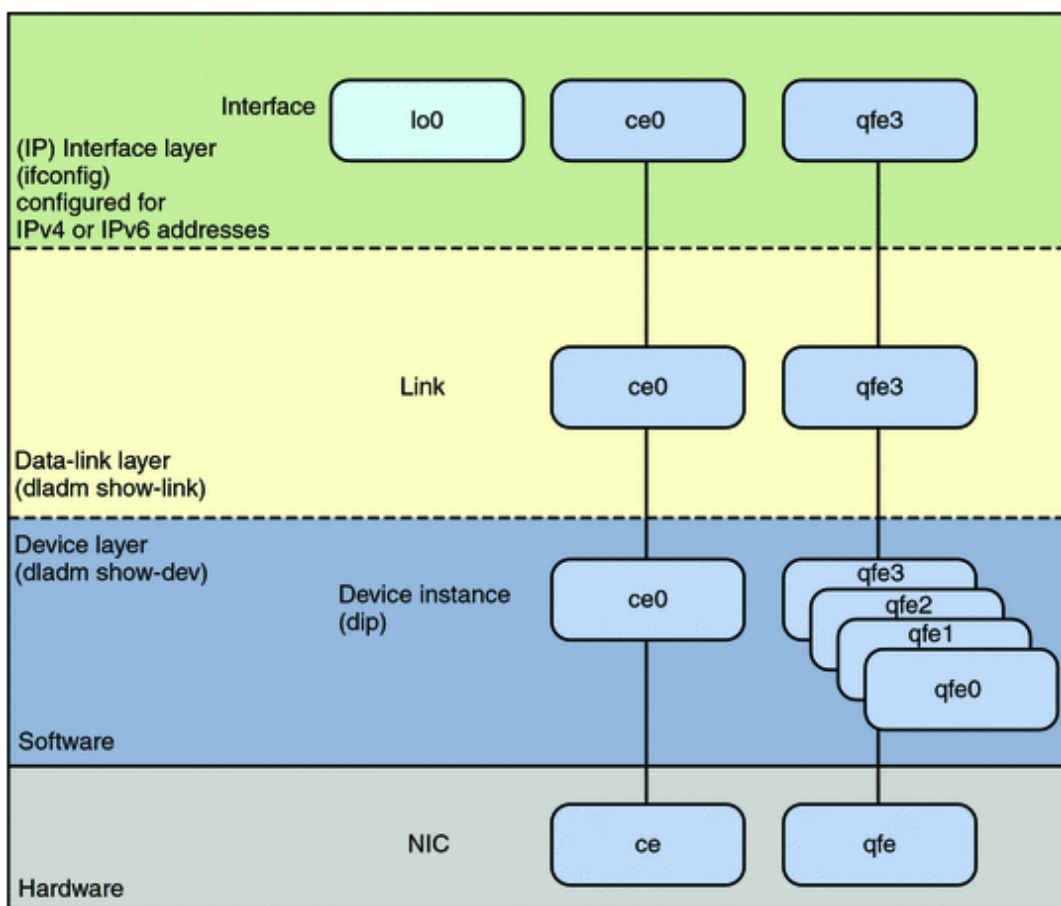
The command `dladm` is used to configure data-link interfaces in Sun Solaris. A configured

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:



Question No : 65 - (Topic 1)

View the Exhibit and review the disk configuration.

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

Primary label contents:

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

format> █
```

The following command is executed on the disk:

```
zpool create pool1 c3t3d0s0
```

What is the result of executing this command?

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

Answer: B

Question No : 66 - (Topic 1)

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

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

Answer: C

Explanation:

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

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

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

Example:

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

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

Question No : 67 - (Topic 1)

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.

Question No : 68 - (Topic 1)

You are troubleshooting a newly installed desktop Oracle Solaris 11 system with a single network interface. From this system, you can connect to other systems within the company intranet, but cannot access any external services (such as websites and email), even when using IP addresses.

Examining the routing table confirms that the default route to 192.168.1.1 is missing. DHCP is not used at this site. Which two commands will temporarily and permanently configure the default route?

- A. `ipadm set-gateway 192.168.1.1`
- B. `route add default 192.168.1.1`
- C. `ipadm set-default 192.168.1.1`
- D. `dladm route-add -d 192.168.1.1`
- E. `echo 192.168.1.1 >/etc/gateway`
- F.

Answer: B,F

Explanation:

B: Setting the default route on Solaris is easy. If you are trying to just set the route temporarily you can use the route command:

```
Route add default <ipaddress>
```

Example:

```
Route add default 192.168.1.1
```

Note: Route command manipulates the kernel routing tables. Routing is the process of forwarding a packet from one computer to another. It is based on the IP address in the IP packet header and netmask.

F: If you want the route to be persisted when you reboot the system, you will need to set the route in the /etc/defaultrouter file.

```
/etc/defaultrouter
```

Example:

```
Echo 192.168.1.1 > /etc/defaultrouter
```

Question No : 69 - (Topic 1)

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 uninstll 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.

Question No : 70 - (Topic 1)

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.

Question No : 71 - (Topic 1)

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

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

Answer: C

Explanation:

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

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

Example:

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

Question No : 72 - (Topic 1)

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: A,B,C,D

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)

Question No : 73 - (Topic 1)

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

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

Answer: E

Explanation:

The ipadm delete-addr subcommand removes addresses from interfaces.

To remove an address from the IPMP group, type the following command:

```
# ipadm delete-addr addrobj
```

The addrobj uses the naming convention inder-interface/user-string.

Question No : 74 - (Topic 1)

Which two SMF milestones can be specified at boot time?

- A. none
- B. network
- C. all
- D. config
- E. unconfig
- F. devices

Answer: A,C

Explanation:

The milestones that can be specified at boot time are

- none
- single-user
- multi-user
- multi-user-server
- all

Question No : 75 - (Topic 1)

In a fresh installation of Oracle Solaris 11, default datalinks are named with a generic naming convention, and they increment as you add interfaces. What is the default name?

- A. eth#
- B. net#
- C. e1000g#
- D. lo#
- E. nic#

Answer: B

Explanation:

When you install this Oracle Solaris release on a system for the first time, Oracle Solaris automatically provides generic link names for all the system's physical network devices. This name assignment uses the net# naming convention, where the # is the instance number. This instance number increments for each device, for example, net0, net1, net2, and so on.

Note:

Network configuration in Oracle Solaris 11 includes

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

Question No : 76 - (Topic 1)

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 use 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

Question No : 77 - (Topic 1)

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.

Question No : 78 - (Topic 1)

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.

Question No : 79 - (Topic 1)