

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 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

NEW QUESTION 3

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

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

201301(Pr201301)

What are the two problems with your plan?

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

Answer: CD

Explanation:

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

D: Group names contain only lowercase characters and numbers.

NEW QUESTION 4

You are going to create live zones on your server. Disk space is critical on this server so you need to reduce the amount of disk space required for these zones. Much of the data required for each of these zones is identical, so you want to eliminate the duplicate copies of data and store only data that is unique to each zone. Which two options provide a solution for eliminating the duplicate copies of data that is common between all of these zones?

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

Answer: DE

Explanation:

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

NEW QUESTION 5

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 6

The core dump configuration in your non-global zone is

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

A user is running a process in a non-global zone (testzone) and the process crashes. The process information is: user126632618017:46:42pts/20:00/usr/bin/bash

When the user's process crashes in testzone, a non-global zone, where will the core dump be saved?

- A. The file will be stored in the non-global zone's directory: /var/core/pprocess/core.hash.2663.
- B. The file will be saved in the global zone's directory: /var/core/core.bash.2663.
- C. A core file cannot be generated in a non-global zone because it shares the kernel with the global zone.
- D. The file will be stored in the global zone's directory: /var/core/pprocess/core.bash.2663.
- E. The file will be saved in non-global zone's directory: /var/core/core.bash.2663

Answer: E

Explanation:

The line
init core file pattern: /var/core/core.%f.%p
will be used for the non-global process to determine the destination of the dump file.
Note: When a process is dumping core, up to three core files can be produced: one in the per-process location, one in the system-wide global location, and, if the process was running in a local (non-global) zone, one in the global location for the zone in which that process was running.

NEW QUESTION 7

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 8

Which files must be edited in order to set up logging of all failed login attempts?

- A. /etc/default/login, /var/adm/loginlog, /etc/syslog.conf
- B. /etc/default/login, /var/adm/authlog, /etc/syslog.conf
- C. /var/adm/loginlog, /var/adm/authlog, /etc/syslog.conf
- D. /etc/default/login, /var/adm/authlog, /var/adm/loginlog

Answer: B

Explanation:

This procedure captures in a syslog file all failed login attempts.

1. Set up the /etc/default/login file with the desired values for SYSLOG and SYSLOG_FAILED_LOGINS
Edit the /etc/default/login file to change the entry. Make sure that SYSLOG=YES is uncommented.
2. Create a file with the correct permissions to hold the logging information. Create the authlog file in the /var/adm directory.
3. Edit the syslog.conf file to log failed password attempts. Send the failures to the authlog file.

NEW QUESTION 9

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

Answer: C

Explanation:

STMF – Manages transactions, such as context and resources for Small Computer System Interface (SCSI) command execution, and tracking logical unit and port providers. STMF also handles logical unit mappings, allocating memory, recovering failed operations, enumeration, and other necessary functions of an I/O stack. STMF is controlled by stmfadm, and stmfadm is the majority of the commands you will be using to administer COMSTAR (COMmon Multiprotocol Scsi TARget).

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

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

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

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

NEW QUESTION 10

You need to update an OS image on a client. The pkg publishers command displays the wrong publisher with the wrong update:

PUBLISHERTYPESTATUSURI

Solaris origin onlinehttp://pkg.oracle.com/solaris/release

The update is available on the updated publisher: PUBLISHERTYPESTATUSURI

Solaris originonlinehttp://sysA.example.com

Select the option that describes the procedure used to update the OS image on the system from the updated publisher.

- A. Copy the repository from the ISO image onto the local client
- B. Configure the repository on the client by using the svccfg - s command so that the Solaris publisher is connected to the new repository
- C. Refresh the application/pkg/server service
- D. Issue the pkgrepo refresh command to refresh the repository catalog
- E. Configure the publisher on the client using the svccfg - s command so that the Solaris publisher is connected to the repository at http://sysA.example.comRefresh the application/pkg/server service
- F. Issue the pkgrepo refresh command to repository catalog
- G. Use the pkg set-publisher command to change the URL of the publisher Solaris to http://sysA.example.co
- H. Issue the pkg update command to update the OS image.
- I. Add the new publisher http://sysA.example.com SolarisUse the pkg set-publisher command to set the publisher search order and place http://sysA.example.com of http://pkg.oracle.com/solaris/releaselIssue the pkg publisher command to view the publisher
- J. Set the new publisher to stick
- K. Issue the pkg update command to update the OS image.

Answer: C

Explanation:

You can use the pkg set-publisher command to change a publisher URI. Changing a Publisher Origin URI

To change the origin URI for a publisher, add the new URI and remove the old URI. Use the -g option to add a new origin URI. Use the -G option to remove the old origin URI.

```
# pkg set-publisher -g http://pkg.example.com/support \  
-G http://pkg.example.com/release example.com
```

Note: You can use either the install or update subcommand to update a package.

The install subcommand installs the package if the package is not already installed in the image. If you want to be sure to update only packages that are already installed, and not install any new packages, then use the update subcommand.

NEW QUESTION 10

Which two SMF milestones can be specified at boot time?

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

Answer: AC

Explanation:

The milestones that can be specified at boot time are none
single-user multi-user
multi-user-server all

NEW QUESTION 12

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 17

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 19

The current ZFS configuration on server is:

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

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

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

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 24

Given the following output of the zpool status command:

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

Identify the correct statement regarding pool1's configuration.

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

Answer: B

NEW QUESTION 26

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

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

Answer: BDFGI

Explanation:

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

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

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

1. The client system boots over the network and gets its network configuration and the location of the install server from the DHCP server.
2. The install server provides a boot image to the client.
3. Characteristics of the client determine which installation instructions and which system configuration instructions are used to install the client.

4. The Oracle Solaris 11 OS is installed on the client, pulling packages from the package repository specified by the installation instructions in the AI install service.
 G: Install the AI tool set.

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

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

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

NEW QUESTION 31

Consider the following rule file for use with the Basic Audit Reporting Tool (BART).

```
CHECK all IGNORE dirmtime
/etc/security
/etc/notices IGNORE contents
/export/home
IGNORE mtime size contents
/var CHECK
```

You are using BART to detect inappropriate changes to the file system. Identify the two correct statements describing the attributes recorded.

- A. /var/dhcp Attribute: size uid gid mode acl
- B. /etc/hosts Attributes: size uid gid mode acl intime dest
- C. /var/spool/mqueue Attribute: size uid gid mode acl dirmtime
- D. /etc/security/exec_attr Attribute: size uid mode acl mtime devnode
- E. /export/home/kate/.profile Attributes: uid gid mode acl dirmtime
- F. /export/home/rick/.profile Attributes: size uid gid mode acl mtime contents

Answer: DF

Explanation:

D: According to line /etc/security F: According to line /export/home

Not E: According to line IGNORE dirmtime

Note: In default mode, the bart compare command, as shown in the following example, checks all the files installed on the system, with the exception of modified directory timestamps (dirmtime):

```
CHECK all IGNORE dirmtime
```

Note 2: The Basic Audit Reporting Tool (BART) feature of Oracle Solaris enables you to comprehensively validate systems by performing file-level checks of a system over time. By creating BART manifests, you can easily and reliably gather information about the components of the software stack that is installed on deployed systems.

BART is a useful tool for integrity management on one system or on a network of systems.

NEW QUESTION 35

You need to install the gzip software package on your system. Which command would you use to find the software package in the configured repository?

- A. pkg search gzip
- B. pkg info gzip
- C. pkg contents gzip
- D. pkginfo gzip
- E. yum list gzip

Answer: A

Explanation:

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

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

NEW QUESTION 40

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

```
export dennie_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 42

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

- A. Use the V2V process to migrate an existing Solaris 8 or 9 non-global zone from a Solaris 10 system to a solaris10 branded zone.
- B. Use the V2V process to migrate an existing Solaris 10 non-global whole root zone from a Solaris 10 system to a solaris10 branded whole root zone.
- C. Install a solaris10 branded zone directly from the Oracle Solaris 10 media.
- D. Migrate an existing 64-bit Solaris 10 system to a solaris10 branded non-global zone using the P2V process.
- E. Use the V2V process to migrate an existing Solaris 10 non-global sparse root zone from a Solaris 10 system to a solaris10 branded sparse root zone.

Answer: BC

Explanation:

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

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

C: How to Install the solaris10 Branded Zone

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

NEW QUESTION 45

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

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

You then enter: `ipadm create-ip net3`

What is the output?

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

Answer: B

Explanation:

According to the exhibit the interface already exists.

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

NEW QUESTION 47

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 51

Which two options describe how to override the default boot behavior of an Oracle Solaris 11 SPARC system to boot the system to the single-user milestone?

- A. from the ok prompt, issue this command: `boot -m milestone=single-user`
- B. From the ok prompt, issue this command: `boot -m milestone/single-user`
- C. From the ok prompt, issue this command: `boot -milestone=single-user`
- D. From the ok prompt
- E. issue this command: `boot -s`
- F. From from the ok prompt, issue this command: `boot -m milestone=s`

Answer: AD

Explanation:

By default, Solaris will boot to the pseudo milestone "all" and start all services. This behaviour can be changed at boot time using either "-s" to reach single-user, or the new SMF option "-m milestone=XXX" (see `kernel(1M)` for a list of the bootable milestones) to select an explicit milestone.

Note: `boot -s` is the same as: `boot -m milestone=single-user`

with the difference being that the former is a lot less to type and is what most SysAdmins will be familiar with.

NEW QUESTION 54

user1, while in his home directory, is attempting to run the following command in his home directory: `cp bigfile verybig`

The system displays the following error:

cp: cannot create verybig: Disc quota exceeded

Your initial troubleshooting shows that the `df -h` command indicates he is at 100% capacity. What command would you use to increase the disk space available to the user?

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

Answer: E

Explanation:

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

NEW QUESTION 59

You are installing the Solaris 11 OE by using the Interactive Text Installer. You have selected the option to automatically configure the primary network controller. Which three items will automatically be configured as a result of this selection?

- A. The IP address.
- B. The name service.
- C. The time zone.
- D. A default user account.
- E. The terminal type.
- F. The root password.
- G. The host name.

Answer: ABC

Explanation:

IP address and name service (such as a DNS server) are provided by the DHCP server.

NEW QUESTION 61

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 deiete-ipv4 ner3/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`.

NEW QUESTION 65

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 68

What is the output of the following command, if executed using the default shell for the root role account of a standard Live CD Install of Oracle Solaris 11?

```
echo '$SHELL'
```

- A. /usr/bin/bash
- B. /usr/bin/ksh
- C. \$SHELL
- D. the PID for the current shell

Answer: C

Explanation:

Single quotes are most strict. They prevent even variable expansion. Double quotes prevent wildcard expansion but allow variable expansion. For example:

```
#!/bin/sh echo $SHELL
```

```
echo "$SHELL"
```

```
echo '$SHELL' This will print:
```

```
/usr/bin/bash
```

```
/usr/bin/bash
```

```
$SHELL
```

NEW QUESTION 71

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

- A. printenv
- B. banner
- C. .version
- D. set-env
- E. show-devs
- F. value version

Answer: BC

Explanation:

B: banner

Displays power-on banner.

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

```
SPARCstation 2, Type 4 Keyboard
```

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

Displays version and date of the boot PROM.

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

NEW QUESTION 73

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. priocntl -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 77

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 82

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 87

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

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

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

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

Answer: AC

Explanation:

A: %PART

The amount of cpu used as a percentage of the total cpu in a processor-set to which the zone is bound. A zone can only have processes bound to multiple processor sets if it is the global zone, or if psrset(1m) psets are used. If multiple binding are found for a zone, it's %PART is the fraction used of all bound psets. For [total] and [system], %PART is the percent used of all cpus on the system.

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

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

NEW QUESTION 89

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

- A)

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

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

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 91

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

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

Answer: B

Explanation:

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

NEW QUESTION 93

User1 is attempting to run the following command: `cp bigfile verybig`

The system displays the following error:

```
cp: cannot create verybig: Disc quota exceeded
```

Your initial troubleshooting shows that the `df -h` command indicates the account is at 100% capacity. What command would you use to determine how much disk space the user has available?

- A. `zfs get quota rpool/export/home/user1`
 B. `zfs userused@user1`
 C. `zfs quota=1M /rpool/export/home/user1`
 D. `df -h | grep user1`

Answer: A

Explanation:

ZFS quotas can be set and displayed by using the `zfs set` and `zfs get` commands. In the following example, a quota of 10 Gbytes is set on `tank/home/bonwick`.

```
# zfs set quota=10G tank/home/bonwick
# zfs get quota tank/home/bonwick NAME PROPERTY VALUE SOURCE
tank/home/bonwick quota 10.0G local
```

NEW QUESTION 98

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

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

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

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

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

Answer: A

Explanation:

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

How to Install the solaris10 Branded Zone

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

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

NEW QUESTION 102

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

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

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

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

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

Answer: A

NEW QUESTION 106

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 109

A change in your company's security policy now requires an audit trail of all administrators assuming the sysadm role, capturing:

There are two commands 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 111

You want to configure an iSCSI target device on your system.

Select the group package required to install this functionality on your system.

- A. storage-server
- B. solaris-small-server
- C. storage-avs
- D. storage-nas

Answer: A

Explanation:

How to Create an iSCSI LUN

The disk volume provided by the server is referred to as the target. When the LUN is associated with an iSCSI target, it can be accessed by an iSCSI initiator.

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

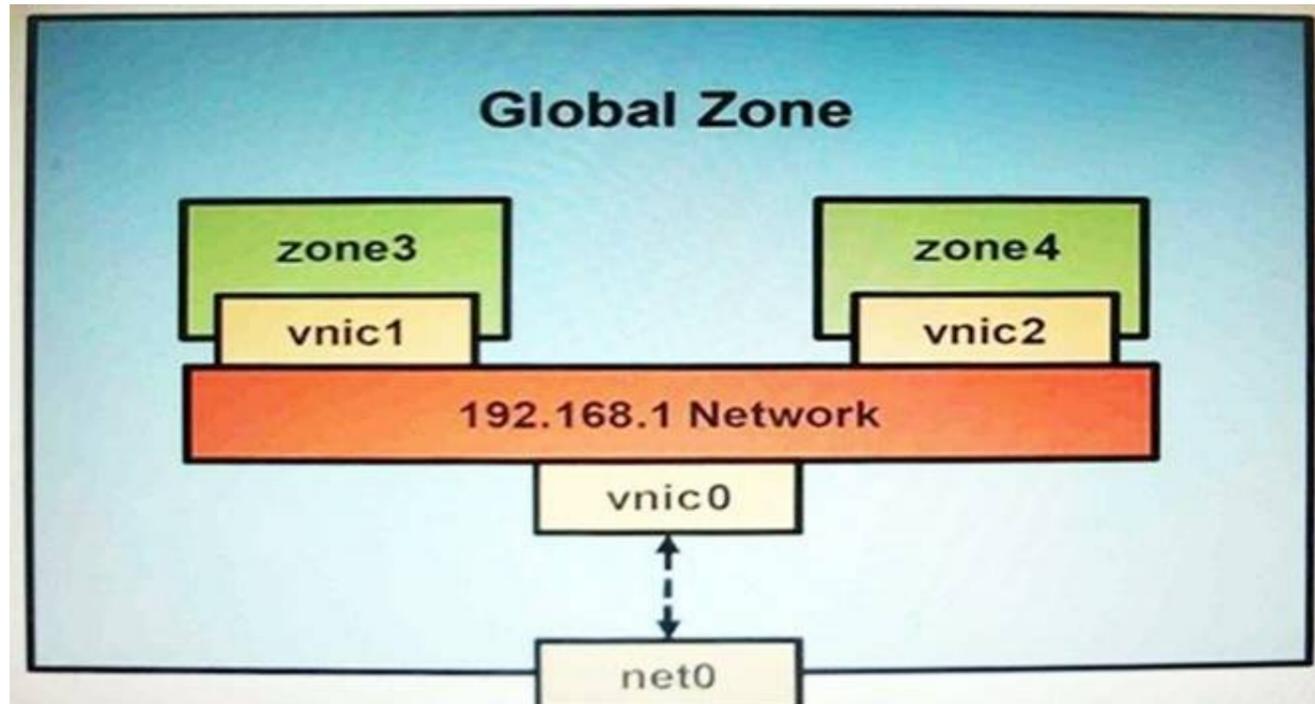
1. Install the COMSTAR storage server software. target# pkg install storage-server

Etc.

NEW QUESTION 114

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 -l vswitch192.168.1 vnic1 fails with the error
 dladm: invalid link name 'vswitch192.168.1' What is the reason for this error?

- A. The name vswitch192.168.1 is not legal.
- B. The zone must be specified withdladm create-vnic -z zone3 vnic1.
- C. The virtual interface must be specified withdladm create-vnic -z zone3 vnic1.
- D. The virtual interface must be created withipadm create-vnic -l switch192.168.1.
- E. The virtual switch must be created first withdladm 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 116

Which operation will fail if the DNS configuration is incorrect?

- A. domainname
- B. ping localhost.
- C. ping 192.168.1.1
- D. ping 23.45.82.174
- E. ping www.oracle.com.
- F. cat /etc/resolv.conf

Answer: E

Explanation:

www.oracle.com would have to be resolved to an IP name by the domain name service.

NEW QUESTION 120

Which two options are characteristics of a fast reboot?

- A. A fast reboot bypasses grub.
- B. A fast reboot cannot be used after a system panic on the x86 platform.
- C. A fast reboot can only be executed on the SPARC platform when the config/fastreboot_default property for the svc:/system/boot-config:default service is set to true.
- D. A fast reboot uses an in-kernel boot loader to load the kernel into memory.
- E. A fast reboot is the default on all platforms.

Answer: CD

Explanation:

C: To change the default behavior of the Fast Reboot feature on the SPARC platform, so that a fast reboot is automatically performed when the system reboots, see below.

The following example shows how to set the property's value to true on the SPARC platform, so that a fast reboot is initiated by default:

```
# svccfg -s "system/boot-config:default" setprop config/fastreboot_default=true
```

```
# svcadm refresh svc:/system/boot-config:default
```

D: Fast Reboot implements an in-kernel boot loader that loads the kernel into memory and then switches to that kernel.

The firmware and boot loader processes are bypassed, which enables the system to reboot within seconds.

The Fast Reboot feature is managed by SMF and implemented through a boot configuration service, svc:/system/boot-config. The boot-config service provides a means for setting or changing the default boot configuration parameters. When the config/fastreboot_default property is set to true, the system performs a fast reboot automatically, without the need to use the reboot -f command. This property's value is set to true on the x86 platform. For task-related information, including how to change the default behavior of Fast Reboot on the SPARC platform, see Accelerating the Reboot Process on an x86 Based System.

Note: One new feature, called Fast Reboot, will allow the system to boot up without doing the routine set of hardware checks, a move that can make system boot times up to two- and-a-half times faster, Oracle claimed. This feature can be handy in that an administrator applying a patch or software update across thousands of Solaris deployments can reboot them all the more quickly.

NEW QUESTION 125

You have a ZFS file system named /dbase/oral and you want to guarantee that 10 GB of storage space is available to that dataset for all data, snapshots, and clones.

Which option would you choose?

- A. zfs set refreservation=10g dbase/oral
- B. zfs set quota=10g dbase/oral
- C. zfs set refquota=10g dbase/oral
- D. zfs set reservation=10g dbase/oral

Answer: D

Explanation:

A ZFS reservation is an allocation of disk space from the pool that is guaranteed to be available to a dataset. As such, you cannot reserve disk space for a dataset if that space is not currently available in the pool. The total amount of all outstanding, unconsumed reservations cannot exceed the amount of unused disk space in the pool. ZFS reservations can be set and displayed by using the zfs set and zfs get commands. For example:

```
# zfs set reservation=5G tank/home/bill
```

```
# zfs get reservation tank/home/bill NAME PROPERTY VALUE SOURCE  
tank/home/bill reservation 5G local
```

NEW QUESTION 130

Which best describes the svc:/system/boot-config service?

- A. It is used to change the milestone on a system.
- B. It is used to set the default run level of the system.
- C. It provides the parameters used to set the system to automatically perform a fast or slow reboot.

D. When the service is enabled, the system performs a fast reboot by default; when it is disabled the system performs a slow reboot by default.

Answer: C

Explanation:

Starting with the Oracle Solaris 11 Express release, Fast Reboot is supported on the SPARC platform, as well as the x86 platform. On both platforms, this feature is controlled by the SMF and implemented through a boot configuration service, svc:/system/boot-config. The boot-config service provides a means for setting or changing the default boot configuration parameters.

The fastreboot_default property of the boot-config service enables an automatic fast reboot of the system when either the reboot or the init 6 command is used. When the config/fastreboot_default property is set to true the system automatically performs a fast reboot, without the need to use the reboot -f command. By default, this property's value is set to false on the SPARC platform and to true on the x86 platform.

NEW QUESTION 133

The line

set noexec_user_stack=1

should be added to the /etc/system file to prevent an executable stack while executing user programs. What is the purpose of this?

- A. help prevent core dumps on program errors
- B. help programs to execute more quickly by keeping to their own memory space
- C. log any messages into the stack log
- D. help make buffer-overflow attacks more difficult

Answer: D

Explanation:

How to Disable Programs From Using Executable Stacks Purpose: Prevent executable stack from overflowing. You must be in the root role.

Edit the /etc/system file, and add the following line: set noexec_user_stack=1

Reboot the system.

reboot

NEW QUESTION 136

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

To accomplish this, you will create .

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

Answer: D

Explanation:

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

NEW QUESTION 139

When setting up Automated Installer (AI) clients, an interactive tool can be used to generate a custom system configuration profile. The profile will specify the time zone, data and time, user and root accounts, and name services used for an AI client installation. This interactive tool will prompt you to enter the client information and an SC profile (XML) will be created.

Which interactive tool can be used to generate this question configuration?

- A. sys-unconfig
- B. installadm set-criteria
- C. sysconfig create-profile
- D. installadm create-profile

Answer: B

Explanation:

Use the installadm set-criteria command to update the client criteria associated with an AI manifest that you already added to a service using installadm add-manifest.

Use the installadm add-manifest command to add a custom AI manifest to an install service.

The value of manifest is a full path and file name with .xml extension. The manifest file contains an AI manifest (installation instructions). The manifest file can also reference or embed an SC manifest (system configuration instructions).

NEW QUESTION 140

Review the boot environment information displayed on your system:

```
oldBE      -   -   149.OK   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 two options accurately describe the newBE boot environment?

- A. It cannot be destroyed.
- B. It cannot be activated.
- C. It cannot be renamed.

- D. You can create a snapshot of it.
- E. It is activated but unbootable.
- F. It has been deleted and will be removed at the next reboot.

Answer: BC

Explanation:

If the boot environment is unbootable, it is marked with an exclamation point (!) in the Active column in the beadm list output.
 The beadm command restricts actions on unbootable boot environments as follows: You cannot activate an unbootable boot environment. (B)
 You cannot destroy a boot environment that is both unbootable and marked as active on reboot.
 You cannot create a snapshot of an unbootable boot environment.
 You cannot use an unbootable boot environment or boot environment snapshot with the -e option of beadm create.
 You cannot rename an unbootable boot environment. (C)

NEW QUESTION 145

Which two statements are true concerning the creation of user accounts by using the useradd command?

- A. By default, it will create the user's home directory.
- B. New user accounts are unlocked but must change their password at their first login.
- C. New user accounts are in a pending activation state until a password is assigned to them.
- D. By default, a new group will be added for each new user account.
- E. By default, the UID of a new user account will be the next available number above the highest number currently assigned.
- F. By default, the UID of a new user account will be the lowest available unused number for nonsystem accounts.

Answer: CE

NEW QUESTION 146

View the following information for a software package:

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

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

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

Answer: B

Explanation:

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

Example:

Listing Information About an Uninstalled Package

```
# pkg info -r php-52 Name: web/php-52 Summary: PHP Server 5.2
Description: PHP Server 5.2 Category: Development/PHP State: Not Installed Publisher: solaris
Version: 5.2.17
Build Release: 5.11
Branch: 0.175.0.0.0.1.530
Packaging Date: Wed Oct 12 14:01:41 2011
Size: 44.47 MB
FMRI: pkg://solaris/web/php-52@5.2.17, 5.11-0.175.0.0.0.1.530:20111012T140141Z
```

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

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

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

NEW QUESTION 151

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

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

F. `pkg update -nv '*'`

Answer: AD

Explanation:

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

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

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

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

NEW QUESTION 153

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 158

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

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

You then ran: `ipadm show-if`

The result indicated that the interface was down.

You then ran:

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

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

The last command indicated that the interface was up.

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

A. The 192.168.0.100 address is reserved for broadcast messages.

B. Another device exists on the network, using the 192.168.0.100 address.

C. The network interface card does not support the address 192.168.0.100.

D. The address 192.168.0.100 is at a boundary and may not be configured in Oracle Solaris 11.

E. 192.168.0.100 is a DHCP address and may not be statically configured in Oracle Solaris 11.

Answer: B

Explanation:

The first IP address is already in use.

NEW QUESTION 159

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

Which command would display the Solaris 11 build version currently running on your system?

A. `pkg info all`

B. `cat /etc/release`

C. `cat /etc/update`

D. `prtconf | grep -i update`

E. `pkg info entire`

Answer: B

Explanation:

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

```
cat /etc/release
```

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

Example:

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

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

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

NEW QUESTION 162

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

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

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

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

Answer: B

Explanation:

Test the remote connection with traceroute.

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

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

NEW QUESTION 166

New features were added to ZFS in Oracle Solaris 11. Your justification to upgrade from Solaris 10 to Oracle Solaris 11 is that it will be possible to take advantage of the enhancements that were made to ZFS.

Identify the three ZFS functions and features that are included in Oracle Solaris 11, but not in Solaris 10.

- A. Encrypted ZFS datasets
- B. Ability for ZFS to detect and remove redundant data from the file system
- C. Shadow Data Migration
- D. Ability to split a mirrored ZFS storage pool
- E. Ability to use ZFS on the boot drive and boot to a ZFS root file system.
- F. elimination of the swap file system when using ZFS on the root disk

Answer: ABC

Explanation:

<http://www.oracle.com/technetwork/server-storage/solaris11/overview/solaris-matrix-1549264.html>

NEW QUESTION 167

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

```
NAME (PUBLISHER)VERSIONINFO
Compress/zip3.1.2-0.175.0.0.0.0.537if-
NAMEVERSIONDATECOMMENT
```

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

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 168

dbzone is currently running on your server.

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

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

Answer: DE

Explanation:

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

E: Use: zlogin zone shutdown

to cleanly shutdown the zone by running the shutdown scripts.

Use this procedure to cleanly shut down a zone.

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

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

NEW QUESTION 169

You attempted to reboot a system via the init command, however the system did not perform boot sequence into the Oracle Solaris Operating Environment. You

are presented with a prompt from the OpenBoot PROM. Which command would you enter, to boot the system from the default device?

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

Answer: B

Explanation:

Boot
 With this form, boot loads and executes the program specified by the default boot arguments from the default boot device
 Note: boot has the following general format: boot [device-specifier] [arguments]
 where device-specifier and arguments are optional.

NEW QUESTION 174

Server A, Server B, and Server C are connected to the same network switch and are on the same network. Each server has a single network interface, net0. You received a tech support call that Server B has lost network connectivity. Your troubleshooting has discovered: Server A can ping Server C, but not Server B. Server B can ping localhost, but not Server A or C. Server C can ping Server A, but not Server B. On Server B, you enter the following command: `dladm show-phys | grep net0`

Response:

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

What is the next logical troubleshooting action?

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

Answer: D

Explanation:

Check the physical connection.

NEW QUESTION 179

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
- 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 184

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

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

Which three statements are correct?

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

Answer: ACD

Explanation:

A: `dir` followed by a single letter.

- C: dir followed by any characters ending with a.
 D: dir followed by any characters, then character b, then one single character. only dirabc matches

NEW QUESTION 185

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: BCD

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.

NEW QUESTION 187

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 190

You are troubleshooting the Oracle Solaris11 Automated Installer (AI), which is not connecting with the IPS software repository.

Which three steps will help determine the cause of DNS name resolution failure?

- A. Verify the contents of /etc/resolve.conf.
- B. Run netstat -nr to verify the routing to the DNS server.
- C. Ping the IP address of the IPS server to verify connectivity.
- D. On the installation server, verify that the menu.1st file for the client points to a valid boot arc hive.
- E. Run df -k to verify that the boot directory containing the boot archive is loopback mounted under /etc/netboot.
- F. Run the command /sbin/dhcpinfo DNSserv to ensure that the DHCP server providing the DNS server information.

Answer: ABF

Explanation:

Check DNS

- * (A) Check whether DNS is configured on your client by verifying that a non-empty /etc/resolv.conf file exists.
- * (F) If /etc/resolv.conf does not exist or is empty, check that your DHCP server is providing DNS server information to the client:
/sbin/dhclient -v
If this command returns nothing, the DHCP server is not set up to provide DNS server information to the client. Contact your DHCP administrator to correct this problem.
- * (B) If an /etc/resolv.conf file exists and is properly configured, check for the following possible problems and contact your system administrator for resolution:
 - ** The DNS server might not be resolving your IPS repository server name.
 - ** No default route to reach the DNS server exists.

NEW QUESTION 194

Examine this command and its output:

```
# zoneadm list Global
zone2
```

Which two outcomes can be deduced from this output?

- A. There is exactly one nonglobal zone installed.
- B. There is one nonglobal zone running.
- C. There is at least one oneglobal zone configured.
- D. There is one nonglobal zone that is not configured.
- E. There is one nonglobal zone that is not running.
- F. There is one nonglobal zone that is not installed.
- G. There is exactly one nonglobal zone configured.

Answer: BC

NEW QUESTION 197

You have installed the SMF notification framework to monitor services. Which command is used to set up the notifications for a particular service?

- A. svccfg
- B. svcadm
- C. setnotify
- D. smtp-notify

Answer: A

Explanation:

How to Set Up Email Notification of SMF Transition Events

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 1:

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
```

Example 2:

The following command creates a notification that sends email when the switch service goes into the online state.

```
# /usr/sbin/svccfg -s svc:/system/name-service/switch:default setnotify to-online \ mailto:sysadmins@example.com
```

Note: The svccfg command manipulates data in the service configuration repository. svccfg can be invoked interactively, with an individual subcommand, or by specifying a command file that contains a series of subcommands.

Changes made to an existing service in the repository typically do not take effect for that service until the next time the service instance is refreshed.

NEW QUESTION 199

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 svcadm command to enable the notification service.
- E. Set the -g maintenance option on the netnotify service to send an email when a service enters the maintenance state.

Answer: C

Explanation:

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

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

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

2. Set notification parameters.

Example:

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

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

NEW QUESTION 201

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:

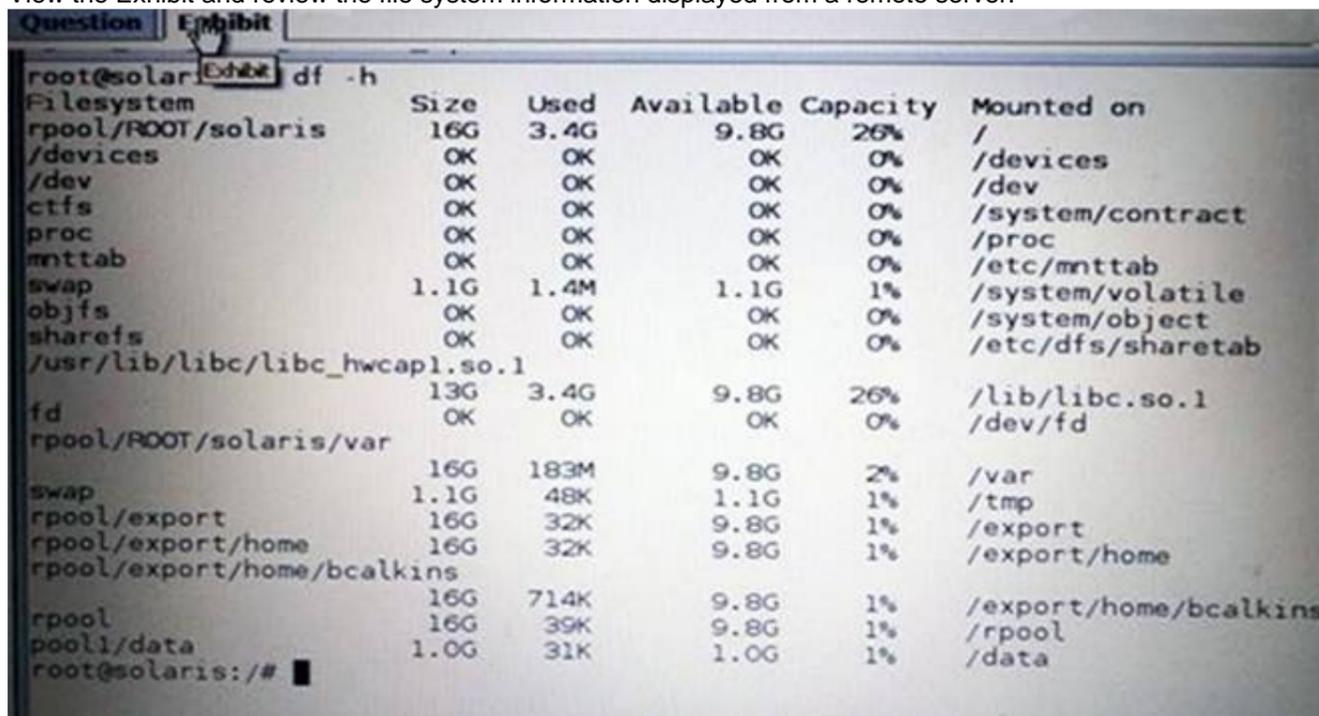
You might need to replace a disk in the root pool for the following reasons: The root pool is too small and you want to replace it with a larger disk. The root pool disk is failing. In a non-redundant pool, if the disk is failing so that the system won't boot, you'll need to boot from an alternate media, 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::disk/c1t0d0
<Physically remove failed disk c1t0d0>
<Physically insert replacement disk c1t0d0>
# cfgadm -c configure c1::disk/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 202

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



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 pool1zfs create pool1/data
- C. zfs set mountpoint=none pool1zfs 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

NEW QUESTION 205

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 207

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



Which two statements describe the Information conveyed by this Icon?

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

Answer: BC

Explanation:

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

C: All online (Wireless)

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

Note that at least one online connection is wireless.

NEW QUESTION 211

Which three files must be edited in order to set up logging of all failed login attempts?

- A. /var/adm/authlog
- B. /etc/syslog.conf
- C. /etc/default/login
- D. /var/adm/loginlog

Answer: ABC

Explanation:

How to Monitor All Failed Login Attempts

This procedure captures in a syslog file all failed login attempts. Assume the Primary Administrator role, or become superuser.

1. Assume the Primary Administrator role, or become superuser.
2. (C) Set up the /etc/default/login file with the desired values for SYSLOG and SYSLOG_FAILED_LOGINS
3. (A) Create a file with the correct permissions to hold the logging information. Create the authlog file in the /var/adm directory.
4. (B) Edit the syslog.conf file to log failed password attempts.

NEW QUESTION 215

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

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

Answer: ACE

Explanation:

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

* If no custom AI manifests are defined for this install service, the default AI manifest is used. The default AI manifest is not associated with any client criteria etc. Each client can use any number of system configuration profiles. If a client system does not use any configuration profile, then an interactive tool opens on that client at first boot after that client installation to complete the configuration of that client.

NEW QUESTION 216

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

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

Answer: C

Explanation:

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

NEW QUESTION 221

Which three statements are true concerning Image Packaging System (IPS) incorporation package?

- A. Installing an incorporation package does not install any other packages.
- B. Every feature or tool has a separate IPS incorporation.
- C. They constrain the versions of packages they incorporate.
- D. They are a content management tool and not a version management tool.
- E. Their dependencies are always of TYPE-REQUIRE.
- F. They are defined by their manifest

Answer: ACE

NEW QUESTION 222

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

```

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

```

The global zone displays the following network information:

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

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

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

Answer: CEG

Explanation:

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

Note on zones:

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

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

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

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

NEW QUESTION 227

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

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

Answer: E

Explanation:

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

name.

NEW QUESTION 232

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 234

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

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

Answer: C

NEW QUESTION 237

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

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

NAME	STATE	READ	WRITE	CKSUM
pool1	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
c3t3d0	ONLINE	0	0	0
c3t4d0	ONLINE	0	0	0
mirror-1	ONLINE	0	0	0
c3t5d0	ONLINE	0	0	0
c3t6d0	ONLINE	0	0	0

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

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

Part	Tag	Flag	First Sector	Size	Last Sector
0	usr	wm	256	145.99GB	306167774
1	unassigned	wm	0	0	0
2	unassigned	wm	0	0	0
3	unassigned	wm	0	0	0
4	unassigned	wm	0	0	0
5	unassigned	wm	0	0	0
6	unassigned	wm	0	0	0
8	reserved	wm	306167775	8.00MB	306184158

Which two are true regarding the ZFS storage pool?

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

Answer: AE

NEW QUESTION 240

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 used. This indicates an automatic network configuration.

NEW QUESTION 242

A user jack, using a korn shell, requests a directory listing as follows:
 jack@solaris:/export/home/jack \$ ls File filea Filea fileb Fileb filec Filec
 Which two statements are correct?

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

Answer: AC

Explanation:

A: starting with one single character, second character must be letter i, any characters, ending with letter a.
 C: starting with letter e, f, or g, followed by anything.

NEW QUESTION 245

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 247

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 252

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 253

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

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

Answer: ABCF

Explanation:

A (not H): You can use your existing storage pools as long as they are upgraded. You have the flexibility of encrypting specific file systems.
 B (not E): Can I enable encryption on an existing pool?

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

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

NEW QUESTION 257

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

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

Answer: B

Explanation:

Kill -9

Kill (terminates without cleanup)

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

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

NEW QUESTION 258

.....

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](#)