

1Z0-821 Dumps

Oracle Solaris 11 System Administrator

<https://www.certleader.com/1Z0-821-dumps.html>



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

The storage pool configuration on your server is:

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

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

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

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

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 3

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

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

Answer: B

Explanation:

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

NEW QUESTION 4

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

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

201301(Pr201301)

What are the two problems with your plan?

- A. Group names may not contain a numeric character

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

Answer: CD

Explanation:

- C: The Group ID (GID) field contains the group's numerical ID. GIDs can be assigned whole numbers between 100 and 60000.
- D: Group names contain only lowercase characters and numbers.

NEW QUESTION 5

You have installed software updates to a new boot environment (BE) and have activated that the booting to the new BE, you notice system errors. You want to boot to the last known good configuration.

Which option would you use on a SPARC system to boot to the currentBE boot environment?

- A. boot -L currentBE
- B. boot -Z rpool/ROOT/currentBE
- C. boot -a Enter the currentBE dataset name when prompted.
- D. boot rpool/ROOT/currentBE
- E. boot -m currentBE
- F. beadm activate currentBE

Answer: F

Explanation:

You can change an inactive boot environment into an active boot environment. Only one boot environment can be active at a time. The newly activated boot environment becomes the default environment upon reboot.

How to Activate an Existing Boot Environment

1. Use the following command to activate an existing, inactive boot environment: beadm activate beName

beName is a variable for the name of the boot environment to be activated. Note the following specifications.

beadm activate beName activates a boot environment by setting the bootable pool property, bootfs, to the value of the ROOT dataset of the boot environment that is being activated.

beadm activate sets the newly activated boot environment as the default in the menu.lst file.

2. Reboot.

The newly activated boot environment is now the default on the x86 GRUB menu or SPARC boot menu.

NEW QUESTION 6

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 7

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

- A. Subnet
- B. unicast
- C. netmask
- D. multicast
- E. broadcast

Answer: C

Explanation:

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

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

NEW QUESTION 8

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 9

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

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

Answer: A

Explanation:

Use the pkg search command to search for packages whose data matches the specified pattern.
Like the pkg contents command, the pkg search command examines the contents of packages. While the pkg contents command returns the contents, the pkg search command returns the names of packages that match the query.

NEW QUESTION 10

You need to connect two nonglobal zones using a private virtual network. Identify the network resources required in the global zone to accomplish this.

- A. an etherstub and two virtual network interfaces
- B. a virtual bridge
- C. two virtual network interfaces.
- D. two etherstubs

Answer: A

NEW QUESTION 10

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 15

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

- A. Updates are only available to customers with an active support contract
- B. The updates are distributed through the My Oracle Support web portal and installed in a central location
- C. All software packages are then updated manually from the command line using the `smpatch` command.
- D. Patches are downloaded from <http://support.oracle.com> either automatically or manually
- E. All software packages are then updated manually from the command line using the `smpatch` or `patchadd` commands.
- F. Software updates are published as packages to a repository
- G. All software packages are then updated manually from the command line using the `pkg` command.
- H. Software updates, published as packages to an OS image
- I. All software packages are then updated manually from the command line using the `pkg` command.

Answer: C

Explanation:

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

```
# pkg update
```

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

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

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

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

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

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

NEW QUESTION 20

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 21

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

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

Answer: C

Explanation:

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

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

```
# ipadm show-addr
```

```
ADDROBJ TYPE STATE ADDR
```

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

NEW QUESTION 25

_____ serves as the interface between the SMF repository and the user to ensure that a consistent picture of the repository is presented to the user.

- A. `repository.db`
- B. `service manifest`
- C. `svc.startd`
- D. `svc.configd`

Answer: D

Explanation:

SVC.CONFIGD is the repository daemon responsible for maintaining /etc/svc/repository.db. The repository.db must come clean during this integrity check otherwise it is a "no go" for usual boot sequence to run level 3. The repository may get corrupted due to various hardware issues, software bugs, disk write failures, etc.

Note: When svc.configd(1M), the Solaris Repository Daemon, is started, it does an integrity check of the smf(5) repository, stored in /etc/svc/repository.db. This integrity check can fail due to a disk failure, the database file being corrupted either due to a hardware bug, a software bug, or an accidental overwrite. If the integrity check fails, svc.configd will write a message to the console.

NEW QUESTION 28

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 31

You want to deploy Oracle Solaris 11 with the Automated Installer (AI). You need to make sure that your server and network meet the requirements for using AI. Identify two requirements for using AI.

- A. You should set up DHC
- B. The DHCP server and AI install server can be the same machine or two different machines.
- C. You can create only one manifest per install servic
- D. If you need more than one manifest, you should create multiple install services.
- E. The minimum requirement to operate as an AI install server is 1 GB of memory.
- F. If two client machines have different architectures and need to be installed with the same version of the Oracle Solaris 11 OS, then you should create two AI manifests and a single install service.

Answer: AD

Explanation:

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

Step 1. The client system boots over the network and gets its network configuration and the location of the install server from the DHCP server.

Step 2: The install server provides a boot image to the client. Etc.

D: If two client machines need to be installed with the same version of the Oracle Solaris 11 OS but need to be installed differently in other ways, then create two AI manifests for the AI install service. The different AI manifests can specify different packages to install or a different slice as the install target, for example.

NEW QUESTION 32

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 36

A datalink can best be described as .

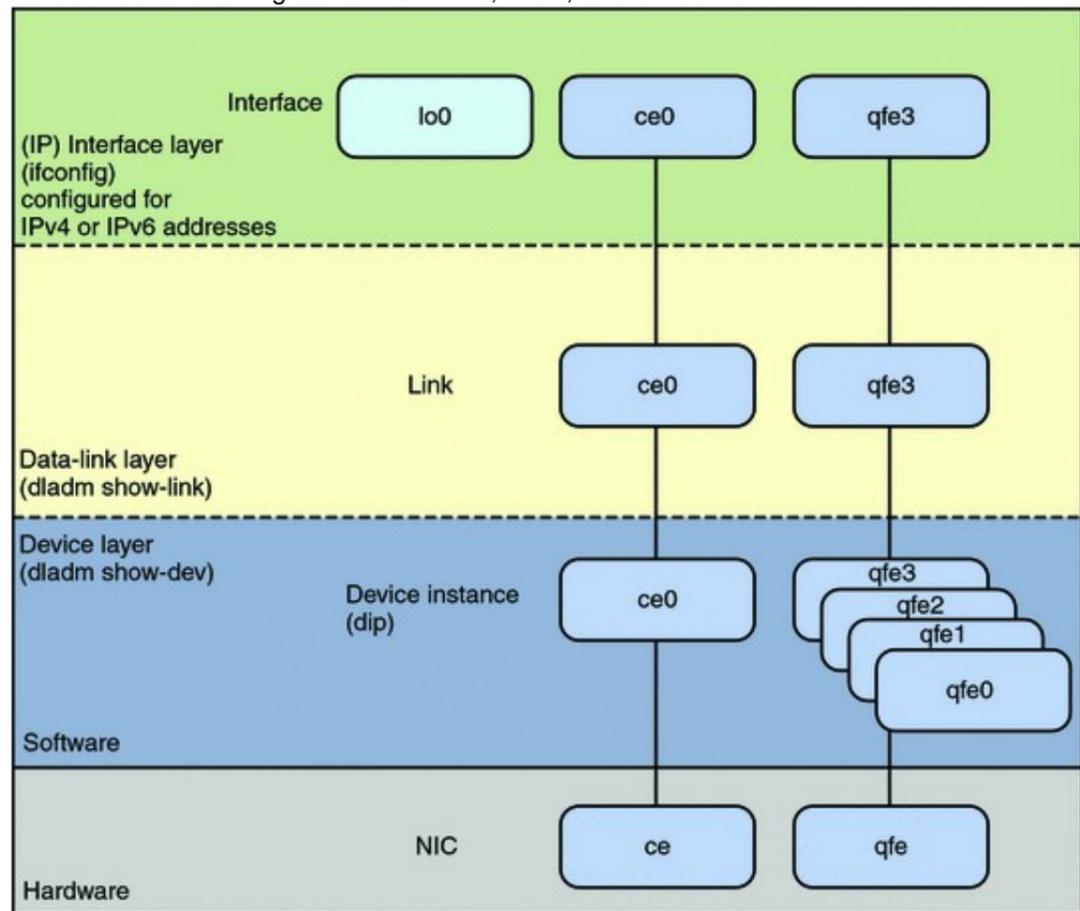
- A. a driver for a Network Interface Card
- B. the software connecting the Internet Layer and the Physical Layer
- C. a device that provides Classless Inter-Domain Routing
- D. a logical object used for IP Multipathing

Answer: D

Explanation:

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

Network Stack Showing Network Devices, Links, and Interfaces:



NEW QUESTION 38

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 39

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 43

In Oracle Solaris 11, where is the Oracle default repository located?

- A. `/var/spool/pkg`

- B. <http://localhost/solaris>
- C. <http://pkg.oracle.com/solaris/release>
- D. <http://www.oracle.com/Solaris/download>
- E. /cdrom/cdrom0

Answer: C

Explanation:

REPOSITORY DESCRIPTION

* <http://pkg.oracle.com/solaris/release>

The default repository for new Oracle Solaris 11 users. This repository receives updates for each new release of Oracle Solaris. Significant bug fixes, security updates, and new software may be provided at any time for users to install at Oracle's discretion.

* <https://pkg.oracle.com/solaris/support>

Provides bug fixes and updates. Accessible with a current support contract from Oracle.

* <https://pkg.oracle.com/solaris/dev> Provides the latest development updates. Accessible to users enrolled in the Oracle Solaris 11 Platinum Customer Program and approved Oracle Partners.

NEW QUESTION 44

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 47

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 51

Oracle Solaris 11 kernel encounters a fatal error, and it results in a system panic.

What type of file does this generate?

- A. a.out
- B. objdump
- C. core dump
- D. tape dump
- E. crash dump

Answer: C

Explanation:

A kernel panic is a type of error that occurs when the core (kernel) of an operating system receives an instruction in an unexpected format or when it fails to handle properly. A kernel panic can also follow when the operating system can't recover from a different type of error. A kernel panic can be caused by damaged or incompatible software or, more rarely, damaged or incompatible hardware.

When a server kernel panics it abruptly halts all normal system operations. Usually, a kernel process named panic() outputs an error message to the console and stores debugging information in nonvolatile memory to be written to a crash log file upon restarting the computer. Saving the memory contents of the core and associated debugging information is called a "core dump."

NEW QUESTION 53

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

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

Answer: B

Explanation:

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

NEW QUESTION 57

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

You have these requirements:

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

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

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

Answer: A

Explanation:

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

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

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

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

config_params/shadow_threads

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

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

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

NEW QUESTION 59

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

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

Answer: B

Explanation:

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

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

NEW QUESTION 61

Which two statements are true concerning the network stack on Oracle Solaris 11?

- A. Hardware network interfaces and datalinks have a one-to-one relationship.
- B. IP addresses are assigned to datalinks.
- C. A single IP interface can have either an IPv4 address or an IPv6 address but not both.
- D. A single IP interface can have both an IPv4 address and an IPv6 address.
- E. A single datalink can have only one IP interface.

Answer: AD

NEW QUESTION 64

The global zone has 8 CPUS. YOU suspect that one of your non global /ones, dbzone, is consuming all of the CPU resources.

Which command would you use to view the CPU utilization for all of the zones to confirm this?

- A. Run from the global zone:prstat -Z
- B. Run from each zone:login <zonename> mpstat
- C. Run from the global zone:zonestar -r summary
- D. Run from the global zone:rctladm -1

E. Run from the global zone:prctl -i

Answer: A

Explanation:

If you're logged on to the system, you can run prstat -Z to generate a summary of cpu/memory utilization by zone.

NEW QUESTION 69

Which network protocol is responsible for routing packets from one network to another?

- A. TCP
- B. UDP
- C. IP
- D. ICMP
- E. Ethernet

Answer: C

Explanation:

The Internet Protocol (IP) is the principal communications protocol in the Internet protocol suite for relaying datagrams across network boundaries. Its routing function enables internetworking, and essentially establishes the Internet.

NEW QUESTION 73

How should you permanently restrict the non-global zone testzone so that it does not use more than 20 CPU shares while it is running?

- A. While configuring the zone, add this entry:add rct1set name = capped.cpu-sharesadd value (priv = privileged, limit = 20, action = none)endexit
- B. While configuring the zone, add this entry: add rct1set name= zone.cpu-sharesadd value (priv=privileged, limit=20, action=none)endexitfrom command line, enter: # dispadmin- d FSS
- C. From the command line enter: #prctl -n zone.cpu-shares - r - v 20 - i zone testzone
- D. From the command line, enter:#prctl - n zone.cpu-shares - v 80 - r - i zone global

Answer: C

Explanation:

The prctl utility allows the examination and modification of the resource controls associated with an active process, task, or project on the system. It allows access to the basic and privileged limits and the current usage on the specified entity.

How to Change the zone.cpu-shares Value in a Zone Dynamically This procedure can be used in the global zone or in a non-global zone.

For more information about roles, see Configuring and Using RBAC (Task Map) in System Administration Guide: Security Services.

```
# prctl -n zone.cpu-shares -r -v value -i zone zonename
```

idtype is either the zonename or the zoneid. value is the new value.

Note: project.cpu-shares

Number of CPU shares granted to a project for use with the fair share scheduler

NEW QUESTION 77

Oracle Solaris 11 limits access to the system with usernames and passwords.

The usernames are held in , and the passwords are held in . Select the correct pair.

- A. /etc/security/policy.conf /etc/passwd
- B. /etc/passwd /etc/shadow
- C. /etc/security /etc/passwd
- D. /etc/shadow /etc/passwd

Answer: B

Explanation:

The /etc/passwd file contains basic user attributes. This is an ASCII file that contains an entry for each user. Each entry defines the basic attributes applied to a user.

/etc/shadow file stores actual password in encrypted format for user's account with additional properties related to user password i.e. it stores secure user account information. All fields are separated by a colon (:) symbol. It contains one entry per line for each user listed in /etc/passwd file.

NEW QUESTION 79

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

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

Answer: D

Explanation:

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

```
# ipadm show-addr
```

```
ADDROBJ TYPE STATE ADDR
```

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

NEW QUESTION 82

alice is a user account used by Alice on a Solaris 11 system. sadmin is a role account on the same system.

Your task is to add the command /usr/sbin/cryptoadm to the Network management profile, so that Alice can execute it, while assuming the sadmin role.

Select the three activities necessary to accomplish this.

- A. To the file /etc/security/prof_attr, add the line: Network Management: solaris:cmd:RO::/usr/sbin/cryptoadm:uid=0
- B. To the file /etc/security/auth_attr, add the line:Network Management:solaris:cmd:RO::/usr/sbin/cryptoadm:uid=0
- C. To the file /etc/security/exec_attr.d/local-entries, add the line:Network Management: solaris:cmd:RO::/usr/sbin/cryptoadm:uid=0
- D. Run the roles alice to ensure that alice may assume the role sadmin.
- E. Run the command profiles sadmin to ensure that the role sadmin includes the network Management profile.
- F. Run the command profiles alice to ensure that the Alice has permissions to access the Network management profile.
- G. Run the command profiles "Network management" to ensure that the Network management profile includes the sadmin role.

Answer: CDG

Explanation:

C: /etc/security/exec_attr is a local database that specifies the execution attributes associated with profiles. The exec_attr file can be used with other sources for execution profiles, including the exec_attr NIS map and NIS+ table.

A profile is a logical grouping of authorizations and commands that is interpreted by a profile shell to form a secure execution environment.

NEW QUESTION 83

Which two statements are true when updating Solaris 11 from one Support Respository Update (SRU) to another SRU by using the pkg update command?

- A. By default, the pkg update command automatically creates a backup Boot Environment whenever the kernel is affected by the update.
- B. By default, the pkg update command automatically creates a new Boot Environment whenever the kernel is affected by the update.
- C. The pkg update command can only be used to update to a newer SRU.
- D. The pkg update command can be used to update to a newer or older SRU.
- E. By default, the pkg update command always updates Solaris 11 to the first SRU that was released after the Current SRU.
- F. The pkg update command can only be performed while running in the single-user milestone.

Answer: BC

NEW QUESTION 84

Given:

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

Which two commands will be successful?

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

Answer: CF

Explanation:

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

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

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

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

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

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

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

NEW QUESTION 89

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

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

Answer: ABD

Explanation:

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

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

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

NEW QUESTION 91

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 94

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

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

Answer: B

Explanation:

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

```
$ suPassword:
```

```
# crontab -l
```

```
#ident "@(#)root 1.19 98/07/06 SMI" /* SVr4.0 1.1.3.1 */
```

```
#
```

```
# The root crontab should be used to perform accounting data collection.
```

```
#
```

```
#
```

```
10 3 * * * /usr/sbin/logadm
```

```
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
```

```
30 3 * * * [ -x /usr/lib/gss/gsscred_clean ] && /usr/lib/gss/gsscred_clean
```

```
#10 3 * * * /usr/lib/krb5/kprop_script slave_kdcs
```

NEW QUESTION 95

What is the result of executing the following command? `svcs -d svc:/network/ssh:default`

- A. disables the `svc:/network/ssh:default` service
- B. displays the services that `svc:/network/ssh:default` is dependent on
- C. displays the services that are dependent on the `svc:/network/ssh:default` service
- D. deletes the `svc:/network/ssh:default` service

Answer: B

Explanation:

The `svcs` command displays information about service instances as recorded in the service configuration repository.

`-d` Lists the services or service instances upon which the given service instances depend.

NEW QUESTION 97

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

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

Answer: D

Explanation:

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

NEW QUESTION 100

Which two accurately describe the Solaris IPS repository?

- A. It contains a collection of operating system patches.

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

Answer: BE

Explanation:

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

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

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

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

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

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

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

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

NEW QUESTION 105

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 108

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 113

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 114

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 117

On server A, you enter the following command to add a static route to serverA route -p add

```
-host 192.168.1.101 192.168.1.101 -static
```

What is the purpose of this command?

- A. to temporarily bypass IP Filter rules
- B. to specify an IPMP target IP address to in.mpathd
- C. to specify routing to an adjacent network when in.rdisc is not used
- D. to specify routing to an adjacent network when in.routed is not used
- E. to ensure the IP address for serverB is not flushed from the ARP cache
- F. to optimize link aggregation using a direct connection between two systems

Answer: B

Explanation:

Note: # route -p add -host destination-IP gateway-IP -static

where destination-IP and gateway-IP are IPv4 addresses of the host to be used as a target.

For example, you would type the following to specify the target system 192.168.10.137, which is on the same subnet as the interfaces in IPMP group itops0:

```
$ route -p add -host 192.168.10.137 192.168.10.137 -static
```

This new route will be automatically configured every time the system is restarted. If you

want to define only a temporary route to a target system for probe-based failure detection, then do not use the -p option.

NEW QUESTION 121

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

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

Which three statements are correct?

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

- D. User jill can edit testfile because she has read and write permission at the group level.
- E. User jack can change permissions for testfile because he is the owner of the file.
- F. User jack can change permissions for testfile because he has execute permission for the file.

Answer: DEF

NEW QUESTION 123

The advantage of core files is that they allow you an opportunity to examine the cause of problems, so that they can be resolved. However, core files must be managed because they .

- A. take up large amounts of disk space
- B. make numerous entries into the /var/adm/wtmpx file
- C. steal resources from the processor, slowing down system performance
- D. fill up swap space; this will begin to slow the system due to swaps
- E. fill up swap space; this will begin to slow the system due to paging

Answer: A

Explanation:

Part of the job of cleaning up heavily loaded file systems involves locating and removing files that have not been used recently. You can locate unused files by using the ls or find commands.

Other ways to conserve disk space include emptying temporary directories such as the directories located in /var/tmp or /var/spool, and deleting core and crash dump files.

Note: Core files are generated when a process or application terminates abnormally. Core files are managed with the coreadm command.

For example, you can use the coreadm command to configure a system so that all process core files are placed in a single system directory. This means it is easier to track problems by examining the core files in a specific directory whenever a process or daemon terminates abnormally.

NEW QUESTION 127

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 129

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

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

Answer: D

Explanation:

pkg/image_root

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

NEW QUESTION 130

You log in to the system as user1, then switch user to root by using the su - command. After entering the correct password, you enter the following commands: whoami;who am i;id

Which option correctly represents the output?

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

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

Answer: B

Explanation:

* The whoami utility displays your effective user ID as a name. Here this would be root.

* who am i

The command who shows who is logged on. Here this would be:

user1 console Dec 30 20:20

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

Here this would be: uid=0(root) gid=0(root)

Note:

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

Real UID

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

Effective UID

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

Saved UID

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

NEW QUESTION 135

You are installing the Oracle Solaris 11 Operating System by using the Text Installer. Which two options describe the features associated with the Text Installer?

- A. It can be used to install only SPARC systems.
- B. It installs gnome as the default user environment on a system capable of displaying a graphical environment.
- C. You can choose whether root is a role or user account.
- D. You can do both automatic and manual configuration of the network.
- E. You can select how to configure the remaining network interfaces.

Answer: CD

NEW QUESTION 139

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

1:07am up 346 day(s), 12:03, 4 users, load average: 0.02, 0.02, 0.01 Which command will be useful in your next step of troubleshooting?

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

Answer: B

Explanation:

Test the remote connection with traceroute.

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

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

NEW QUESTION 141

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

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

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

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

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

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

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

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

Answer: C

NEW QUESTION 143

Review the zonestat command:

```
zonestate -q physical -memory -R high -z -p -p "zones" 10 24h 60m
```

Select the option that correctly describes the information that is displayed by this command.

- A. It is a sample of dbzone's physical memory usage taken every hour over a 24-hour period. Only the top 10 samplings of peak memory usage are displayed.
B. All other utilization data is eliminated.
C. It is a sample of dbzone's CPU, virtual memory, and networking utilization. Physical memory is excluded from the report. The sampling is taken every 10 minutes over a 24-hour period and peak utilization is displayed each hour.
D. It is a sample of dbzone's CPU, virtual memory, and networking utilization. Physical memory is excluded from the report. The sampling is taken every 10 minutes over a 24-hour period and displayed each hour.
E. It is a sample of dbzone's physical memory usage taken every 10 seconds over a 24-hour period. Only peak virtual memory usage and CPU utilization are displayed each hour. All other utilization data is eliminated.
F. It is a sample of dbzone's physical memory usage taken every 10 seconds over a 24-hour period. Only peak memory usage is displayed each hour. All other utilization data is eliminated.

Answer: D

Explanation:

* (Not A, B, C): interval (here 10 seconds): Specifies the length in seconds to pause between each interval report.

* duration (here 24 h)

* -R report[, report] (here high) Print a summary report.

High Print a summary report detailing the highest usage of each resource and zone during any interval of the zonestat utility invocation.

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.

The default output is a summary of cpu, physical, and virtual memory utilization. The -r option can be used to choose detailed output for specific resources.

NEW QUESTION 146

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

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

Answer: ABDFH

Explanation:

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

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

```
$ sudo pkg install slim_install
```

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

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

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

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

NEW QUESTION 147

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

A. cd ~ and cd -

B. cd and cd.

C. cd ~ and cd

D. cd * and cd . .

Answer: A

Explanation:

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

```
cd -
```

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

Note:

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

NEW QUESTION 150

You are creating a non-global zone on your system.

Which option assigns a zpool to a non-global zone, and gives the zone administrator permission to create zfs file system in that zpool?

A. While creating the non-global zone, make the following entry: add deviceset match=/dev/rdisk/c4t0d0endBoot the zone and log in the zone as root

B. Create the zpool: zpool create pool2 c4t0d0In the non-global zone, root can now create ZFS file system in the pool2 zpool

C. In the global zone, create the zpool: global# zpool create pool2 c4t1d0While creating the no-global zone, make the following entry: add datasetset name=pool2endadd fsset dir=pool1set special=pool1set type=zfspool1endBoot the zone, log in the zone as root, and create the zfs file system in the pool2 zpool.

D. In the global zone, create the zpool:global#zpool create pool2 c4t1d0While creating the global zone, make the following entry: add datasetset name=pool2endBoot the zone, log in to the zone as root and create the zfs file systems in the pool2 zpool.

E. In the global zone, create the zpool and the ZFS file systems that you want to use in the non-global zone: global#zpool create pool2 c4t1d0global#zfs create pool2/dataWhile creating the non-global zone, make the following entry for each ZFS file system that you want to make available in the zone: add fsset dir=/dataset special=pool2/dataset type=zfsend

F. Create the zpool in the global zone: global#zpool create pool2 c4t1d0Boot the non- global zone, log in to the zone as root, and issue this command to delegate ZFS permissions to root: non-global zone# zfs allow root create , destroy, mount pool2Log in to the non-global zone create ZFS file systems in the pool2 zpool.

Answer: C

Explanation:

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

NEW QUESTION 155

United States of America export laws include restrictions on cryptography.

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

A. Corporations must utilize signed X.509 v3 certificates.

B. A third-party provider object must be signed with a certificate issued by Oracle.

C. Loadable kernel software modules must register using the Cryptographic Framework SPI.

D. Third-party providers must utilize X.509 v3 certificates signed by trusted Root Certification Authorities.

E. Systems destined for embargoed countries utilize loadable kernel software modules that restrict encryption to 64 bit keys.

Answer: BC

Explanation:

B: Binary Signatures for Third-Party Software

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

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

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

NEW QUESTION 160

You have a ticket from a new user on the system, indicating that he cannot log in to his account. The information in the ticket gives you both the username and password. The ticket also shows that the account was set up three days ago.

As root, you switch users to this account with the following command: su - newuser

You do not get an error message.

You then run `ls -la` and see the following files:

```
local1.cshrc local1.login local1.profile .bash_history .bashrc .profile
```

As root, you grep the `/etc/passwd` file and the `/etc/shadow` file for this username, with these results:

```
/etc/passwd contains newuser:x:60012:10:/home/newuser:/usr/bin/bash
```

```
/etc/shadow contains newuser:UP: : : :10: : As root, what is your next logical step?
```

- A. `Usermod -f 0`
- B. `passwd newuser`
- C. `mkdir /home/newuser`
- D. `useradd -D`

Answer: B

Explanation:

The content of the `/etc/shadow` document indicates that the `newuser` account has no password.

We need to add a password.

The `passwd` utility is used to update user's authentication token(s). D: Here the user account already exist. There is no need to create it.

When invoked without the `-D` option, the `useradd` command creates a new user account using the values specified on the command line plus the default values from the system. Depending on command line options, the `useradd` command will update system files and may also create the new user's home directory and copy initial files.

NEW QUESTION 164

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

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

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

Answer: DE

Explanation:

`FLAGS`

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

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

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

Note:

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

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

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

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

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

NEW QUESTION 169

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 173

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 175

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	Description
*	* 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 177

Which statement is correct about shutdown and init commands?

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

Answer: A

NEW QUESTION 178

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

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

Answer: A

NEW QUESTION 181

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/dhccpinfo 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/dhccpinfo DNSserv`

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 183

Solaris 11 includes a redesigned software packaging model: the Image Packaging system.

Which three describe advantages of the Image Packaging System over the previous Solaris 10 SVR4 packaging model?

- A. Eliminates patching of the software package
- B. Makes the patching process more efficient with less downtime
- C. Eliminates OS version upgrade
- D. Allows for the installation of the OS without a local DVD or installation server
- E. Allows the use of a repository mirror to speed up package operation
- F. Allows users to publish their own software package in a software repository

Answer: AEF

NEW QUESTION 186

Which three Installation option allow for a "hands free" and "unattended" Installation of the Solaris 11 environment?

- A. Jumpstart
- B. LiveCD
- C. A text Installation over the network
- D. An Automated Installation performed on an x86 client
- E. An Automated Installation using media from a local DVD or USB drive
- F. An Automated Installation using a networked repository

Answer: DEF

Explanation:

Oracle Solaris 11 uses Automated Installer (AI) for unattended installations.

Unattended installations are possible by placing the contents of the AI Image media (or ISO image contents from a download) on an AI server.

NEW QUESTION 188

The current ZFS configuration on your server is:

`pool1 124K 3.91G 32K /pool1 pool1/data 31K 3.91G 31K /data`

You need to create a new file system named `/data2`. `/data2` will be a copy of the `/data` file system.

You need to conserve disk space on this server whenever possible.

Which option should you choose to create `/data2`, which will be a read writeable copy of the `/data` file system, while minimizing the amount of total disk space used in `pool1`?

- A. `zfs set mountpoint=/data2 compression=on pool1/data2`
- B. `zfs snapshot pool1/data@nowzfs set mountpoint=/data2, compression=on pool1/data@now`
- C. `zfs create snapshot pool1/data@nowzfs send pool1/data@now | zfs recv pool1/data2`
- D. `zfs create snapshot pool1/data@nowzfs clone -o mountpoint=/data2 pool1/data@now pool1/data2`
- E. `zfs snapshot pool1/data@nowzfs clone -o mountpoint=/data2 -ocompression=on pool1/data@now pool1/data2`
- F. `zfs snapshot pool1/data@nowzfs clone -o mountpoint=/data2 pool1/data@now pool1/data2`

Answer: E

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

`zfs clone [-p] [-o property=value] ... snapshot filesystem|volume` Creates a clone of the given snapshot.

Note:

Because snapshots are fast and low overhead, they can be used extensively without great concern for system performance or disk use .

With ZFS you can not only create snapshot but create a clone of a snapshot.

A clone is a writable volume or file system whose initial contents are the same as the dataset from which it was created. As with snapshots, creating a clone is nearly instantaneous, and initially consumes no additional disk space. In addition, you can snapshot a clone.

A clone is a writable volume or file system whose initial contents are the same as the original dataset. As with snapshots, creating a clone is nearly instantaneous, and initially consumes no additional space.

Clones can only be created from a snapshot. When a snapshot is cloned, it creates an implicit dependency between the parent and child.

NEW QUESTION 192

When you issue the "gzip: zommand not found" message is displayed. You need to install the gzip utility on your system.

Which command would you use to check if the gzip utility is available from the default publisher for installation?

- A. pkg info|grep gzip
- B. pkg list SUNWgzip
- C. pkg contents gzip
- D. pkg search gzip

Answer: D

Explanation:

Searching for Packages

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.

pkg search

search [-H|a|l|p|r] [-o attribute ...] [-s repo_uri] query Search for matches to the query, and display the results.

Which tokens are indexed are action-dependent, but may include content hashes and pathnames.

Note: pkg is the retrieval client for the image packaging system. With a valid configuration, pkg can be invoked to create locations for packages to be installed, called 'images', and install packages into those images. Packages are published by publishers, who may make their packages available at one or more repositories. pkg, then, retrieves packages from a publisher's repository and installs them into an image.

NEW QUESTION 194

Review the storage pool information:

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

Choose the correct procedure to repair this storage pool.

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

Answer: C

Explanation:

You might need to replace a disk in the root pool for the following reasons: The root pool is too small and you want to replace it with a larger disk

The root pool disk is failing. In a non-redundant pool, if the disk is failing so that the system won't boot, you'll need to boot from an alternate media, such as a CD or the network, before you replace the root pool disk.

In a mirrored root pool configuration, you might be able to attempt a disk replacement without having to boot from alternate media. You can replace a failed disk by using the zpool replace command.

Some hardware requires that you offline and unconfigure a disk before attempting the zpool replace operation to replace a failed disk.

For example:

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

NEW QUESTION 198

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

```

root@solaris:~# df -h
Filesystem      Size  Used Available Capacity  Mounted on
rpool/ROOT/solaris 16G  3.4G   9.8G    26%      /
/devices        OK    OK      OK      0%      /devices
/dev            OK    OK      OK      0%      /dev
ctfs            OK    OK      OK      0%      /system/contract
proc           OK    OK      OK      0%      /proc
mnttab         OK    OK      OK      0%      /etc/mnttab
swap           1.1G  1.4M   1.1G     1%      /system/volatile
objfs          OK    OK      OK      0%      /system/object
sharefs        OK    OK      OK      0%      /etc/dfs/sharetab
/usr/lib/libc/libc_hwcapi.so.1
fd             13G  3.4G   9.8G    26%      /lib/libc.so.1
rpool/ROOT/solaris/var
swap           16G  183M   9.8G     2%      /var
swap           1.1G  48K   1.1G     1%      /tmp
rpool/export    16G  32K   9.8G     1%      /export
rpool/export/home 16G  32K   9.8G     1%      /export/home
rpool/export/home/bcalkins
rpool          16G  714K   9.8G     1%      /export/home/bcalkins
pool1/data     16G  39K   9.8G     1%      /rpool
pool1/data     1.0G  31K   1.0G     1%      /data
root@solaris:~#

```

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

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

This new server also has the following file systems configured:

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

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

- A. zfs create -o mountpoint=/data -o refquota=1g pool1/data
- B. zfs set mountpoint=none 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 200

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

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

Answer: C

Explanation:

zonecfg info

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

Note: zonecfg -z

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

Incorrect Answer

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

list option:

Display the name of the current zones, or the specified zone if indicated. B: No such command.

D: no such options zonecfg -icv

NEW QUESTION 201

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

- A. EDITOR=vi
- B. crontab=vi

- C. crontab – e vi
- D. env

Answer: A

Explanation:

Set the EDITOR variable to vi.

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

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

Incorrect answers

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

NEW QUESTION 204

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 208

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 212

Subnets are created by using .

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

Answer: B

Explanation:

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

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

NEW QUESTION 213

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.

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

Answer: A

Explanation:

rctldm - display and/or modify global state of system resource controls
The following command activates system logging of all violations of task.max-lwps.
rctldm -e syslog task.max-lwps
#

NEW QUESTION 223

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 225

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

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

Answer: B

Explanation:

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

NEW QUESTION 229

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

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

Remote system name: backupSYS

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

Most recent backup name: rpool-1202

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

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

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

Answer: A

Explanation:

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

* ZFS root pool cannot be recovered

* ZFS root pool snapshots are stored on a remote system and are shared over NFS

* The system is booted from an equivalent Solaris release to the root pool version so that the Solaris release and the pool version match. Otherwise, you will need to add the -o version=version-number property option and value when you recreate the root pool in step 4 below.

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

Boot from CD/DVD or the network.

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

ok boot net -s

ok boot cdrom -s

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

2.

Mount the remote snapshot dataset. For example:

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

3.

Recreate the root pool. For example:

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

4.

Restore the root pool snapshots.

This step might take some time. For example:

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

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

5.

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

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

Install the boot blocks on the new disk.

On a SPARC based system:

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

NEW QUESTION 230

You want to display the IP address assignments of the network interfaces. Which command should you use?

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

Answer: B

Explanation:

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

```
# ipadm show-addr
```

```
ADDROBJ TYPE STATE ADDR
```

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

NEW QUESTION 232

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 235

In an effort to reduce storage space on your server, you would like to eliminate duplicate copies of data in your server's ZFS file systems.

How do you specify that pool1/data should not contain duplicate data blocks (redundant data) on write operations?

- A. zfs create -o compression=on pool1/data
- B. zpool create -o deduplication=on pool1; zfs create pool1/data
- C. zfs create -o deduplication=on pool1; zfs create pool1/data
- D. zfs create -o dedupratio=2 pool1/data
- E. zfs create -o dedup=on pool1/data

Answer: E

Explanation:

ZFS Deduplication Property

Solaris Express Community Edition, build 129: In this Solaris release, you can use the deduplication 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.

You can enable this property as follows:

zfs set dedup=on tank/home

NEW QUESTION 240

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

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

Answer: C

Explanation:

Pgrep takes a process name and return a PID.

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

Incorrec answers:

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

NEW QUESTION 243

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

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

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

Time Zone: US/Eastern

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

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

Terminal type: sun-color

Users:
No user account

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

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

Answer: BD

NEW QUESTION 246

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

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

Answer: E

Explanation:

The following options are supported:

C

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

Note:

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

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

NEW QUESTION 247

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 249

You are setting up an automated installer (AI) install server and issue the following command:

```
installadm create-service -n prod_ai -s /repo/prod_ai.iso \  
-i 192.168.1.100 -c 5 -d /export/repo
```

Which four options describe the install server that you have configured?

- A. The service name is prod_ai.
- B. DHCP base IP address is 192.168.1.100
- C. The initial IP address for the install clients will be 192.168.1.100. This IP address is temporary.
- D. After the client is booted, it will use IP addresses in the following range: 192.168.1.101-105.
- E. Five IP addresses are allocated for DHCP clients, starting with 192.168.1.100.
- F. The Install server will support up to five clients.
- G. The AI net image ISO file is located in /repo/prod and the net image ISO will be unpacked in /export/repo.
- H. The AI net image ISO file is located in /repo/repo and is named /repo/prod/_ai.iso.

Answer: ABDF

Explanation:

A: -n <svcname>

Uses this install service name instead of default service name.

B: -i <dhcp_ip_start>

Sets up a new DHCP server. The IP addresses, starting from dhcp_address_start, are set up.

D: -c <count_of_ipaddr>

Sets up a total number of IP addresses in the DHCP table equal to the value of the count_of_ipaddr. The first IP address is the value of dhcp_ip_start that is provided by the -i option.

F: -s <srcimage>

Specifies location of AI ISO image to use for setting up the install service.

<targetdir>

Required: Specifies location to set up net image.

NEW QUESTION 254

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

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

Answer: AB

Explanation:

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

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

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

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

Incorrect Answer

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

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

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

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

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

NEW QUESTION 255

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

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

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

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

Answer: A

Explanation:

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

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

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

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

NEW QUESTION 256

Select two statements that correctly describe the capabilities of the Distribution Constructor.

- A. ISO images for use with the Automated Installer (AI) can be created.
- B. Bootable USB images can be created for SPARC and x86 architectures.
- C. A single installation server can be used to create ISO images for SPARC and x86 architectures.
- D. Checkpoints can be used to pause the build, allowing scripts to run that modify the resulting ISO image.
- E. A single installation server can be used to create ISO images for Solaris 10 and Solaris 11 operating systems.

Answer: AD

Explanation:

A: You can use the distribution constructor to create the following types of Oracle Solaris images:

* (A) x86 or SPARC ISO Image for Automated Installations

* Oracle Solaris x86 live CD image

* x86 or SPARC Oracle Solaris text installer image

* x86 Oracle Solaris Virtual Machine

Note: You can use the distribution constructor to build custom Oracle Solaris images. Then, you can use the images to install the Oracle Solaris software on individual systems or multiple systems. You can, also, use the distribution constructor to create Virtual Machine (VM) images that run the Oracle Solaris operating system.

D: Checkpointing Options

You can use the options provided in the `distro_const` command to stop and restart the build process at various stages in the image-generation process, in order to check and debug the image that is being built. This process of stopping and restarting during the build process is called checkpointing.

NEW QUESTION 261

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 1Z0-821 Exam with Our Prep Materials Via below:

<https://www.certleader.com/1Z0-821-dumps.html>