PassQueen
a.k.a dumpsqueen.com

Oracle Solaris 11 System Administration

Oracle 1z0-821

Version Demo

Total Demo Questions: 33

Total Premium Questions: 330

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Managing Software Packages	49
Topic 2, Managing Boot and Shutdown	19
Topic 3, Managing Services	21
Topic 4, Administering User Accounts	28
Topic 5, Managing System Resources	14
Topic 6, Administering ZFS	51
Topic 7, Administering Boot Environments	10
Topic 8, Administering Zones	35
Topic 9, Administering Network	51
Topic 10, Administering Virtual Networks	5
Topic 11, Administering System Security	20
Topic 12, Analyzing System Performance	12
Topic 13, Mix Questions	15
Total	330

QUESTION NO: 1

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

A. pkg search

B. pkg info

C. pkginfo

D. pkg contents -t depend -o fmri solaris-desktop

ANSWER: D

Explanation:

`pkg contents -t depend -o fmri solaris-desktop` is the appropriate command because an IPS group package identifies its member packages through dependency actions in its package manifest. The `pkg contents` subcommand displays manifest actions for an installed package or, when appropriate repository access is available, for a package available from configured publishers. Restricting the displayed action type with `-t depend` limits the output to dependencies, and selecting the `fmri` attribute with `-o fmri` produces the package FMRIs that constitute the software group. This gives an administrator the actual set of package names represented by `solaris-desktop`, which is the information needed before installing or assessing that group. The IPS `pkg` manual documents `contents` as the command for displaying package manifest contents and supports filtering actions by type and selecting output attributes. See the Oracle Solaris pkg(8) manual page for the command syntax and manifest-query behavior. Oracle's Installing and Updating Software documentation also describes IPS group packages and their dependency-based composition.

QUESTION NO: 2

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: A B

Explanation:

"DHCP address assignment cannot be configured in a shared IP zone." is correct because a shared-IP zone uses the global zone's IP-layer configuration and state. The global zone administrator assigns and manages the zone's network address configuration; a shared-IP zone does not run its own independently administered IP stack and cannot use DHCP to obtain its address. This design provides network isolation at the address level while retaining centralized control of interfaces, routing, and IP protocol state in the global zone.

"Shared IP is the default type of network configuration." is also correct. When a zone is configured without explicitly selecting exclusive-IP networking, Solaris uses shared-IP networking by default. In this model, the non-global zone shares the global zone's TCP/IP stack and is normally given one or more addresses on a datalink administered from the global zone. Administrators select exclusive-IP instead only when the zone requires its own IP-related state and independent network-stack administration. Oracle documents the shared-IP default and the distinct management model for each zone networking type in the Oracle Solaris Zones configuration overview and the zonecfg(8) manual page.

QUESTION NO: 3

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

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

A. Concatenate `so1.repo.iso-a` and `so1.repo.iso-b` into `so1.full.iso`, mount the resulting ISO image, and use `rsync` to copy its repository contents to `/export/IPS/repo`. Set `pkg/inst_root` to `/export/IPS/repo` and `pkg/readonly` to `true`, then configure the publisher with `pkg set-publisher -G http://pkg.oracle.com/solaris/release/ -g http://serverA.example.com/ solaris`.

B. cat so1.repo.iso-a so1.repo.iso-b > /export/IPS/repoSet the pkg/inst_root property to true and the pkg/readonly property to /export/IPSSet the preferred publisher by using pkg set-publisher -G http://serverA.example.com/ \-g http://pkg/oracle.com/solaris/rekease/solaris

C. cat so1.repo.iso-a so1.repo.iso-b > so1.full.isoMount the ISO image and use the rsync command to extract the contents of the ISO file to /export/IPS/repoSet the pkg/inst_root property to /export/IPS/repo and the pkg/readonly property to trueSet the preferred publisher by using pkg set-publisher solaris \-g http://pkg.oracle.com/

D. cat so1.repo, iso-a so1.repo.iso-b > /export/IPS/repo.isoMount the ISO image and copy the repo directory from the ISO image to /export/IPS/reposet the pkg/inst_root property and the pkg/readonly property to /export/IPS/reposet the preferred pkg/inst_root property by using pkg set-publisher - G http://serverA.example.com/ \- g http://pkg.oracle.

ANSWER: A

Explanation:

The correct procedure is to concatenate the split ISO files in their proper sequence, producing one complete ISO image before mounting it. The repository content must then be copied from that mounted image into `/export/IPS/repo`; using `rsync` is appropriate because it recursively preserves the repository directory hierarchy and attributes. The IPS depot service is configured with `pkg/inst_root` set to `/export/IPS/repo`, which identifies the repository root it will serve, and with `pkg/readonly` set to `true`, appropriate for a repository copied from installation media. After setting these service properties, the package depot service should be refreshed or restarted so it serves that repository location.

The publisher configuration uses the publisher name `solaris`. The `-g` argument adds `http://serverA.example.com/` as the origin from which packages are obtained, while `-G` removes the specified Oracle origin so the local repository is used as the configured source. This produces a local, read-only IPS repository that clients can access through ServerA. Oracle documents publisher-origin management in the pkg(8) manual page and package depot configuration in the pkg.depotd(8) manual page.

QUESTION NO: 4

You have already generated a 256-bit AES raw key and named the keystore file /mykey. You need to use the key to create an encrypted file system.

Which command should you use to create a ZFS encrypted file system named pool1/encrypt using the /mykey keystore?

A. zfs create - o encryption = /mykey pool1/encrypt

B. zfs create -o encryption=aes-256-ccm -o keysource=raw,file:///mykey pool1/encrypt

C. zfs create - o encryption = AES keysource = /mykey pool1/encrypt

D. zfs create - o encryption = on keystore = /mykey pool1/encrypt

ANSWER: B

Explanation:

`zfs create -o encryption=aes-256-ccm -o keysource=raw,file:///mykey pool1/encrypt` is correct because it creates the requested ZFS file system and supplies both encryption configuration properties at creation time. The `encryption` property selects the AES-256-CCM encryption algorithm, which is the appropriate Solaris ZFS property value for a 256-bit AES key in this scenario. The `keysource` property identifies how ZFS obtains the key. Its `raw` type tells ZFS that the source contains key material in raw form rather than a passphrase or another key format. The URI `file:///mykey` names the local file that holds the already generated raw key. The comma between `raw` and the file URI is part of the required `keysource` property syntax, while each `-o` introduces a separate ZFS property. Finally, `pool1/encrypt` is the dataset name to create. Oracle's Solaris documentation illustrates this same pattern: generate a 256-bit AES raw key, then create the encrypted dataset with `encryption=aes-256-ccm` and a `keysource=raw,file:///...` setting. See the Solaris zfs(1M) manual page and Oracle's Solaris documentation portal.

QUESTION NO: 5

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

To accomplish this, you will create_____.

A. an ether stub

B. virtual router

C. a virtual bridge

D. a virtual network interface

E. nothing, because a virtual switch is automatically created when the virtual network interfaces are created

ANSWER: A

Explanation:

An ether stub is the required foundation for this type of isolated, host-local virtual network in Oracle Solaris. An ether stub is a software-defined Ethernet switch that has no connection to a physical network interface. Virtual network interfaces can be created on top of that ether stub and assigned to the local zones that need to communicate. Because the ether stub is not backed by a physical link, frames exchanged by those virtual network interfaces remain within the Solaris system and are separated from traffic using other links, switches, or zones.

After creating the ether stub, the administrator creates a virtual network interface for each participating zone, specifying the ether stub as that interface's underlying link. The zones can then use normal IP configuration over those interfaces, while the ether stub provides their dedicated Layer 2 segment. This model is useful for private application tiers, internal service paths, and test networks that must not be exposed to the host's external network. Oracle documents ether stubs as virtual switches used to connect virtual network interfaces in an isolated virtual network. See the Oracle Solaris virtual network documentation and the Oracle Solaris network virtualization administration guide.

QUESTION NO: 6

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. boot -m milestone=single-user

B. boot -m milestone/single-user

C. boot -milestone=single-user

D. boot -s

E. boot -m milestone=s

ANSWER: A D

Explanation:

Oracle Solaris 11 uses Service Management Facility milestones to control the level of services started during boot. The normal default is the `all` milestone, whereas the `single-user` milestone starts only the limited services needed for system administration and recovery. At the SPARC OpenBoot PROM prompt, `boot -m milestone=single-user` passes an SMF milestone override to the kernel, explicitly instructing it to boot to the single-user milestone. The `boot -s` form is the documented shorthand for this same single-user boot behavior and remains a standard recovery and maintenance boot command for Solaris administrators. Both commands are temporary boot-time overrides: they affect the current boot rather than permanently changing the system's configured default milestone. This is useful when maintenance is required, when diagnosing startup problems, or when a full multiuser service environment should not be started. Oracle documents the `-m milestone=` boot argument and identifies `-s` as the single-user equivalent in the Solaris boot documentation. See the kernel(1M) manual page and Oracle's booting and shutting down guidance.

QUESTION NO: 7

You have Solaris 11 system with a host name of sysA and it uses LDAP as a naming service.

You have created a flash archive of sysA and you want to migrate this system to an Oracle Solaris11 server, Solaris10 branded zone.

The zone Status on the Oracle Solaris 11 server is: - zone10 incomplete/zone/zone1solaris10exc1

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

A. Use zonecfg to change the zonename before booting the system for the first time

B. Use the - u option with the zoneadm - z zone10 attach command.

C. Use the -u option with the zoneadn -z zone10 install command.

D. Remove the sysidcfg file from the /root directory before booting the non-global zone.

E. Use the -u option with the zoneadm -z zone10 install command.

ANSWER: E

Explanation:

Use the `-u` option with the `zoneadm -z zone10 install` command. For a Solaris 10 branded zone installed from a Flash archive, `-u` unconfigures the zone as part of installation. This removes the source system's inherited system-configuration identity, including information such as its hostname and configured naming-service details. Consequently, when the installed non-global zone is first booted, Solaris runs the system-configuration process and prompts for the required identity and network naming-service values. This is particularly important when the archive originated from `sysA`, because the migrated zone must not silently retain that host's LDAP-related configuration and hostname. The option belongs on the installation operation used to deploy the archive, for example `zoneadm -z zone10 install -u -a /path/to/archive.flar`. Oracle documents `zoneadm install -u` as installing the zone in an unconfigured state so that configuration is performed at initial boot. See the Oracle zoneadm(1M) manual page and the Oracle Solaris 10 branded zones administration documentation.

QUESTION NO: 8

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: B

Explanation:

The software connecting the Internet Layer and the Physical Layer is the best description of a datalink in the context of the Solaris networking stack. A datalink is a Layer 2 networking abstraction that provides the connection between IP-layer interfaces and the underlying network hardware or virtual networking resource. It is the object through which frames can be sent and received over a network medium, using link-layer characteristics such as MAC addressing. Solaris treats this layer as a distinct administrative object, allowing it to represent not only a physical network interface card but also software-defined networking constructs, including link aggregations, VLANs, virtual NICs, and etherstubs.

IP interfaces are configured on top of datalinks, while the datalink itself is associated with the physical device or virtual resource that provides actual network connectivity. The `dladm` utility administers these link-layer objects, whereas IP-related configuration is managed separately. This separation is a core Solaris networking design feature because it enables flexible configuration without making the IP interface synonymous with the hardware device. Oracle documents datalinks as the Layer 2 entities managed by `dladm` in the Solaris network virtualization documentation and in the dladm(8) manual page.

QUESTION NO: 9

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

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

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

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

A. Option A

B. Option B

C. Option C

D. Option D

E. Option E

F. Option F

G. Option G

H. svcadm clear svc:/network/ssh:default

ANSWER: H

Explanation:

The minimum recovery command is `svcadm clear svc:/network/ssh:default`. In Solaris Service Management Facility, a service instance that has entered the `maintenance` state is intentionally held there after its method has failed repeatedly or an administrator has placed it in maintenance. Clearing the instance removes that administrative maintenance condition and directs SMF to reevaluate the instance, satisfy its dependencies, and retry the configured start method. If the underlying transient condition has already been resolved, the SSH service transitions through its normal states and returns online without requiring a separate disable/enable sequence or a system reboot.

Using the full FMRI, `svc:/network/ssh:default`, precisely identifies the default SSH service instance and avoids ambiguity if similarly named services or instances exist. The `clear` subcommand is specifically intended for this maintenance-recovery workflow; it resets the service's maintenance state so SMF can again manage and start it according to the service manifest and repository configuration. Oracle documents the `svcadm clear` operation in the svcadm(8) manual page and describes SMF service state administration in the Solaris SMF documentation.

QUESTION NO: 10

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: B D E

Explanation:

User jack can use `cat` to output the contents of `testfile` because he has read permission as the file owner. Solaris evaluates file access using the applicable permission class. Since jack owns `testfile`, the owner permissions shown for the file govern his ordinary read access, and read permission permits utilities such as `cat` to read its contents.

User jill can edit `testfile` because she has read and write permission at the group level. Her membership in the file's group causes the group permission class to apply to her. Read access allows the editor to read the existing contents, while write access allows it to save changes to the regular file. This is the normal permission combination required to edit an existing file.

User jack can change permissions for `testfile` because he is the owner of the file. On Solaris, the owner of a file may use `chmod` to modify its mode bits; this authority comes from file ownership rather than from any particular read, write, or execute bit currently set on the file. Oracle's Solaris documentation describes the owner/group/other permission model and the owner's ability to alter modes with `chmod`: Solaris chmod(1) manual page and Managing File and Directory Permissions.

QUESTION NO: 11

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

Which command is best?

A. kill -2 1234

B. kill -15 1234

C. kill -9 1234

D. kill -1 1234

ANSWER: B

Explanation:

`kill -15 1234` is the best command because signal 15 is `SIGTERM`, the standard termination-request signal. An orderly shutdown means giving the target process an opportunity to handle termination cleanly: it can catch the signal, stop accepting new work, finish or roll back active operations as appropriate, close files and network connections, remove

temporary resources, and write any required state before exiting. This is the normal signal used when an administrator wants a process to terminate gracefully rather than be immediately forced out of memory.

In fact, `SIGTERM` is also the default signal sent when `kill` is invoked without a signal argument, so `kill 1234` would have the same termination-request behavior. Specifying `-15` explicitly makes the intended signal clear and is appropriate for the stated requirement. Signal names or numbers may be supplied to `kill`; the POSIX specification identifies `TERM` as the signal used for normal termination requests. Solaris administrators should generally begin process shutdown with `SIGTERM`, allowing application-level cleanup handlers to run before considering more forceful action.

See the POSIX kill utility specification and Oracle's Solaris kill(1) manual page.

QUESTION NO: 12

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: A E F

Explanation:

Image Packaging System delivers software as versioned packages from package repositories, rather than relying on the separate SVR4 package and patch workflow. Therefore, *Eliminates patching of the software package* is correct in the IPS sense: fixes and enhancements are delivered as updated package versions, and administrators use package update operations to move an image to the appropriate package versions. IPS also supports repository mirrors. *Allows the use of a repository mirror to speed up package operation* is an advantage because a nearby or local mirror can provide the same package content with lower latency and reduced dependency on a remote repository. Finally, IPS is designed around publishers and repositories, so *Allows users to publish their own software package in a software repository* is correct. An organization can create IPS packages, publish them to its own repository, and configure systems to consume them through a publisher. This supports controlled internal delivery of locally developed or third-party software alongside Oracle-provided packages. Oracle documents IPS package repositories, publishers, package update behavior, and package publication tools in its Solaris packaging documentation: Oracle Solaris 11.2 Packaging and Delivering Software With IPS and Configuring Publishers and Repositories.

QUESTION NO: 13

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: A B C F

Explanation:

Oracle Solaris ZFS native encryption is available for newly created ZFS datasets, including file systems and volumes, after the pool has the required on-disk feature level. For a pool created under Solaris 10, upgrading the pool to ZFS pool version 30 enables the encryption capability; an encrypted file system can then be created in that upgraded pool. The existing unencrypted file system is not converted in place, so data intended for protection is copied or migrated to the newly created encrypted dataset.

Encryption is built into the standard ZFS administration model. Administrators use the normal `zfs` command set and ZFS properties to create encrypted datasets and manage their keys, without installing a separate “ZFS Dataset Encryption” package. This integration also permits encryption to be selected for individual datasets rather than requiring every file system to use the same protection arrangement. In practical terms, the pool can contain both encrypted and unencrypted datasets, while encryption settings and keys are managed in the context of the encrypted dataset.

Oracle’s Solaris documentation describes ZFS encryption requirements, dataset-level encryption configuration, and key management in the Managing ZFS Encryption section. Pool upgrade behavior is covered in Oracle’s Upgrading ZFS Pools documentation.

QUESTION NO: 14

On which is the open boot prom available?

A. x86 only

B. x86 64-Bit only

C. SPARC only

D. both x86 and x86 64-Bit

E. x86, x86 64-Bit and SPARC

ANSWER: C

Explanation:

OpenBoot PROM is available on SPARC systems. OpenBoot is the firmware environment used by SPARC-based Oracle systems to perform early hardware initialization, provide device and boot-management commands, and load the Solaris boot program. Its interactive command interface is commonly reached at the `ok` prompt, where an administrator can inspect devices, select boot targets, alter boot arguments, and perform diagnostic or recovery-related actions before Solaris starts. In Solaris administration, commands such as `boot`, `probe-scsi`, and `printenv` are associated with this SPARC OpenBoot environment. Oracle’s Solaris documentation distinguishes the SPARC OpenBoot firmware boot process from the x86 boot process, which uses PC firmware and a boot loader rather than OpenBoot PROM. Consequently, SPARC is the applicable platform for OpenBoot PROM in the context of Solaris 11 administration. Oracle provides a description of OpenBoot and its role in SPARC system startup in the Oracle Solaris Administration: Devices and File Systems guide. Further OpenBoot command and firmware reference material is available in Oracle’s OpenBoot documentation.

QUESTION NO: 15

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 gid mode acl mtime contents

E. /export/home/kate/.profile Attributes: uid gid mode acl dirmtime

F. /export/home/rick/.profile Attributes: size uid gid mode acl mtime

ANSWER: D F

Explanation:

The corrected attribute set for `/etc/security/exec_attr` is `size uid gid mode acl mtime contents`. The initial `CHECK all` directive enables recording of every BART-checkable attribute. The subsequent `IGNORE dirmtime` removes only the directory modification-time attribute. Because `exec_attr` is a regular file below `/etc/security`, its applicable checked attributes include its size, ownership, permissions, ACL, modification time, and file contents. Directory-specific attributes such as `dirmtime`, and device-specific attributes such as `devnode`, do not belong in this file's recorded set.

The corrected attribute set for `/export/home/rick/.profile` is `size uid gid mode acl mtime`. The rule in effect for `/export/home` retains the prior exclusion of `dirmtime` and adds an exclusion for `contents`. Therefore, BART records the regular file's metadata and modification time, but does not calculate or record its content attribute. This is useful for home-directory trees when administrators want to detect ownership, permission, size, and timestamp changes without treating changes to users' file contents as integrity events.

Oracle documents BART rule processing and the supported check attributes in the bart(1M) manual page and the bart_rules(7) manual page.

QUESTION NO: 16

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. After the client is booted, it will use IP addresses in the following range: 192.168.1.101-105.

D. Five IP addresses are allocated for DHCP clients, starting with 192.168.1.100.

E. The Install server will support up to five clients.

F. The AI source ISO file is located at /repo/prod_ai.iso, and the net image will be unpacked in /export/repo.

G. The AI net image ISO file is located in /repo/repo and is named /repo/prod/_ai.iso.

ANSWER: A B D F

Explanation:

The service name is `prod_ai` because the `-n` option explicitly assigns that name to the newly created Automated Installer service. The DHCP base IP address is `192.168.1.100`, as supplied by `-i`; this is the first address in the DHCP address range configured for installation clients. Five IP addresses are allocated for DHCP clients, starting with `192.168.1.100`, because `-c 5` specifies the count of addresses beginning at the address provided through `-i`. Thus, the DHCP allocation comprises five sequential client addresses starting at that base address.

The AI source ISO is `/repo/prod_ai.iso`, as specified by `-s`. The `-d /export/repo` argument identifies the target directory in which the AI network image is established from that source image. Together, these arguments create a named AI install service, populate its network boot/install image in the target location, and configure the associated DHCP address allocation required to boot AI clients. Oracle documents these `create-service` options in the installadm(1M) manual page and the AI service workflow in the Oracle Solaris Automated Installer documentation.

QUESTION NO: 17

You start to execute a program by using the following command:

~/bigscript &

You then determine that the process is not behaving as expected, and decide that you need to terminate the process.

Based on the information shown below, what is the process number you should terminate?

```
#echo $$
15156
# ps -aef | grep 15156
    root 15163   15156  0  12:51:15  pts/3    0:00 bash
    root 15156    5420  0  12:33:15  pts/3    0:00 bash
    root 15166   15156  0  12:51:45  pts/3    0:00 grep
    root 15165   15156  0  12:51:45  pts/3    0:00 ps -aef
```

A. 15163

B. 15156

C. 15166

D. 15165

ANSWER: A

Explanation:

The process number to terminate is **15163**. Appending `&` to `~/bigscript` causes the shell to start the script as a background job and then immediately return a prompt. In the displayed process hierarchy, process 15156 is the interactive shell, while 15165 and 15166 are the temporary `ps` and `grep` commands used to inspect the process list. The remaining child process, 15163, is therefore the background shell/script process created for `~/bigscript`.

Sending a termination signal to 15163 stops the background job without terminating the user's login shell or merely stopping the diagnostic commands. In Solaris, a process can be signaled by its process ID using `kill`; for example, `kill 15163` requests normal termination with `SIGTERM`. If the script has started additional child processes, process-group handling may also be relevant in practice, but the PID identified by this process listing as the launched background script is 15163. Oracle documents the process information reported by `ps` in the ps(1) Solaris manual page and signal delivery by PID in the kill(1) Solaris manual page.

QUESTION NO: 18

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: D

Explanation:

No response; the command was successful. is correct because `dladm show-phys` displays physical data links, while `ipadm create-ip` creates an IP interface on top of an existing data link. Thus, the presence of `net3` in the physical-link display establishes that it is an available link on which an IP interface can be created; it does not itself mean that an IP interface named `net3` has already been created.

The syntax `ipadm create-ip net3` is valid for creating an IPv4/IPv6 IP interface associated with the `net3` datalink. No IP address object is required at this stage. Address configuration is a separate operation, such as using `ipadm create-addr` after the IP interface exists. On successful completion, this administrative command normally produces no terminal output. The resulting IP interface can subsequently be verified with commands such as `ipadm show-if`.

Oracle documents `create-ip` as the subcommand used to create an IP interface over a datalink in the ipadm(1M) manual page. The distinction between data links administered by `dladm` and IP interfaces administered by `ipadm` is also described in the Oracle Solaris network administration documentation.

QUESTION NO: 19

Which three options accurately describe Oracle Solaris 11 zones?

A. can be NFS servers

B. are whole root type only

C. cannot have their own time zone setting

D. can execute z£s and zpool commands (from a non-global zone)

E. are virtualized operating system environments, each with its own dedicated OS and kernel

F. are virtualized operating system environments, created with a single instance of the OS shared kernel

G. can have their own time zone setting

ANSWER: A F G

Explanation:

Oracle Solaris zones share one running Solaris kernel while isolating applications, processes, users, networking, and administrative configuration into separate operating system environments. Therefore, “are virtualized operating system environments, created with a single instance of the OS shared kernel” accurately states the fundamental Solaris Zones architecture. The global zone owns the kernel and physical hardware, while non-global zones run in isolated user-space environments.

“Can be NFS servers” is also correct. Solaris supports providing network services from non-global zones, including NFS service, allowing an application or file-serving workload to be isolated from other workloads on the same system. This isolation is a central operational use of zones.

“Can have their own time zone setting” is correct because a zone’s configuration can specify a distinct time zone. Although non-global zones normally use the system clock maintained by the global zone and cannot independently change that clock, their local time-zone presentation can differ through the zone configuration. Oracle documents zones as isolated environments sharing a single kernel and documents configurable zone properties, including time-zone behavior. See the Oracle Solaris Zones introduction and the zonecfg(8) manual page.

QUESTION NO: 20

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

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

A. svccfg -s svc:/system/system-log setprop start/exec= "syslogd -t"
Restart the syslogd daemon.

B. Set the following parameter in the /etc/syslogd.conf file: LOG_FROM_REMOTE=YES
Restart the syslogd daemon.

C. svcadm enable svc:/system/system-log/config/log_from_remote
Restart the syslogd daemon.

D. svccfg -s svc:/system/system-log setprop config/log_from_remote=true
Restart the syslogd daemon.

E. Set the following parameter in the /etc/default/syslogd file: LOG_FROM_REMOTE=YES
Restart the syslogd daemon.

ANSWER: D E

Explanation:

The supported Solaris 11 service-management method is to set the `config/log_from_remote` Boolean property on the `svc:/system/system-log` service to `true`, then restart the service. This property controls whether the service starts `syslogd` configured to receive and log messages submitted by remote hosts. SMF is the preferred persistent configuration mechanism because it keeps the setting under service management and applies it whenever the service is started.

Setting `LOG_FROM_REMOTE=YES` in `/etc/default/syslogd` and restarting the daemon is also a valid configuration mechanism. The `LOG_FROM_REMOTE` setting determines whether messages from remote systems are logged; enabling it

allows syslogd to accept remote log traffic. After changing either persistent configuration setting, restart the system log service so that the running daemon receives the new configuration. Remote clients must, of course, be able to reach the system's syslog service over the relevant network path.

Oracle documents the system-log SMF service and its properties in the syslogd(8) manual page and documents SMF property administration in the svccfg(8) manual page.

QUESTION NO: 21

Select two correct statements about the authentication services available in Oracle Solaris 11.

A. Pluggable Authentication Modules (PAM) is used to control the operation of services such console logins and ftp.

B. The Secure Shell can be configured to allow logins across a network to remote servers without transmitting passwords across the network.

C. Secure Remote Procedure Calls (Secure RPC) provides a mechanism to encrypt data on any IP Socket connection.

D. Pluggable Authentication Modules (PAM) is used to implement the Secure Shell in Oracle Solaris 11.

E. Simple Authentication and Security Layer (SASL) provides a mechanism to authenticate and encrypt access to local file system data.

ANSWER: A B

Explanation:

Pluggable Authentication Modules (PAM) is used to control the operation of services such console logins and ftp. Oracle Solaris uses PAM as a configurable authentication framework for applications and services, including login-related services. PAM service configuration specifies the modules and policies used when a user attempts to authenticate, establish an account session, or perform related access-control operations. This design lets administrators apply authentication policy to individual services without requiring those services to be rewritten. Oracle documents PAM configuration in the `pam.conf` and `pam.d` service-policy files; see the pam(7) Oracle Solaris manual page.

The Secure Shell can be configured to allow logins across a network to remote servers without transmitting passwords across the network. In particular, SSH public-key authentication proves possession of a user's private key by signing authentication data, rather than sending the account password to the remote host. SSH also protects the connection with encryption, making it the appropriate remote-login service where credentials and session traffic must be protected from network interception. Oracle Solaris documents public-key user authentication and the required key files in its ssh(1) manual page.

QUESTION NO: 22

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:

The command `pkg info -r gzip` displays detailed, human-readable package metadata for the `gzip` package when it is not installed in the current image. The `pkg info` subcommand reports package attributes such as the package name, summary, description, category, publisher, version, build release, packaging date, size, and FMRI. Ordinarily, `pkg info` obtains this information from packages installed locally. The `-r` option changes the source to the package repositories associated with the image's configured publishers. This permits IPS to retrieve metadata for available packages and report their state as not installed, provided that a package pattern such as `gzip` is supplied. The system therefore needs access to an appropriate configured publisher repository, but the package itself does not need to be installed. This is the suitable command when the desired output is the package-information display rather than merely a package name, file list, or integrity result. Oracle documents that `pkg info -r` retrieves package information from configured publisher repositories and requires one or more package patterns. See the pkg(1) manual page and Oracle's Displaying Package Information documentation.

QUESTION NO: 23

You want to install the openldap software package to a now boot environment for testing before introducing the now software package to the production environment. What option describes the correct procedure to:

1) create a new BE named nowBE

2) install the software to that new BE only

A. pkg install --newBE openldap

B. pkg install --be-name nowBE openldap

C. beadm create newBEbeadm mount newBE /mntpkg -R /mnt update openldap

D. beadm create newBEbeadm activate newBEpkg install openldap

ANSWER: B

Explanation:

`pkg install --be-name nowBE openldap` is the correct procedure because IPS can create a new boot environment as part of a package operation and install the requested package into that newly created environment. The `--be-name` option supplies the name for the new boot environment, so `nowBE` is created as a clone of the active boot environment before the OpenLDAP package is changed. The package installation is then applied to the clone, leaving the currently active production boot environment unmodified. This is particularly useful for testing: after installation, the administrator can activate `nowBE` with `beadm activate nowBE` and reboot into it when ready to validate the change. If testing is unsuccessful, the existing boot environment remains available for booting without needing to undo the package operation. Oracle documents that package operations may create a new boot environment and that `--be-name` specifies its name. See the Oracle Solaris pkg(1) manual page and Oracle's Boot Environments documentation.

QUESTION NO: 24

dbzone is currently running on your server.

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

A. zlogin –z dbzone halt

B. zoneadm –z dbzone shutdown –i0

C. zoneadm –z dbzone shutdown

D. zoneadm –z dbzone halt

E. zlogin dbzone shutdown –i0

ANSWER: B C E

Explanation:

The `zoneadm -z dbzone shutdown -i0`, `zoneadm -z dbzone shutdown`, and `zlogin dbzone shutdown -i0` commands all request an orderly shutdown from within the zone's normal operating environment. A zone shutdown uses the zone's init/SMF shutdown processing, allowing services and applications to receive their normal stop requests and execute their configured shutdown procedures. Specifying `-i0` requests transition to init state 0, the powered-off state; when omitted from `zoneadm shutdown`, the normal shutdown behavior still performs an orderly transition of the running zone to the installed state. Running `shutdown` through `zlogin` similarly executes the shutdown utility inside `dbzone`, so the zone's own service-management framework can stop managed services cleanly before the zone is brought down. Oracle documents `zoneadm shutdown` as the administrative operation for shutting down a running zone and documents `zlogin` as the mechanism for executing a command in a non-global zone. See the zoneadm(1M) manual page and the zlogin(1) manual page.

QUESTION NO: 25

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

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

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

A. dbzone is using 0.21% of the total CPU resource available in the zone's processor set.

B. dbzone is using 0.21% of the global zone's total CPU.

C. dbzone is using 5.48% of the total physical memory that has been allocated to the zone.

D. dbzone is using 2.37% of the global zone's total virtual memory.

E. The network is being utilized 100% with no physical bandwidth remaining.

ANSWER: A C

Explanation:

The correct interpretations are "dbzone is using 0.21% of the total CPU resource available in the zone's processor set." and "dbzone is using 5.48% of the total physical memory that has been allocated to the zone." The `zonestat` report provides both system-relative utilization and utilization relative to the limits or resource partitions applicable to an individual zone. In the CPU section, the `%PART` value expresses CPU consumption as a percentage of the processor-set capacity to which the zone is bound. Therefore, the displayed value of 0.21 represents dbzone's use of the CPU resources in its applicable processor-set partition, rather than a general statement about CPU assigned to another zone. In the physical-memory

section, the percentage associated with the zone's capacity indicates the portion of the physical-memory allocation or cap currently consumed by dbzone. The displayed 5.48 value consequently represents the proportion of physical memory allocated to that zone which is in use. Oracle documents that `zonestat` reports zone CPU, memory, network, and resource-control utilization at specified intervals, including percentages of configured limits and system resources. See the Oracle zonestat(1) manual page and the Oracle Solaris Zones documentation.

QUESTION NO: 26

Select the packet type that identifies members of the group and sends information to all the network interfaces in that group.

A. Unicast

B. Multicast

C. Broadcast

D. Bayesian

E. Quality of Service Priority

ANSWER: B

Explanation:

Multicast is the correct packet delivery type because it is designed for one-to-many communication with an explicitly defined group of receiving interfaces. An IPv6 multicast address identifies a group of interfaces, which may reside on one or more nodes. When a sender transmits an IPv6 packet addressed to that multicast group, the network delivers a copy to every interface that has joined the group. This allows applications and network services to efficiently distribute the same information to multiple interested recipients without separately sending an individual unicast packet to each recipient.

In Solaris 11 IPv6 networking, multicast is fundamental to several protocol functions, including neighbor discovery and router discovery. IPv6 does not use broadcast addresses; multicast addresses serve the group-delivery role where earlier protocols might have used broadcasting. IPv6 multicast addresses begin with the `ff00::/8` prefix, and their scope controls how broadly packets can be propagated, such as link-local or organization-local delivery. Oracle's Solaris documentation describes IPv6 multicast as addressing a group of interfaces and delivering packets to all members of that group. See Oracle Solaris IPv6 overview and the IPv6 Addressing Architecture (RFC 4291).

QUESTION NO: 27

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

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

Select the three activities necessary to accomplish this.

A. To the file /etc/security/prof_attr, add the line: Network Management: solaris:cmd:RO::/usr/sbin/cryptoadm:euid=0

B. To the file /etc/security/auth_attr, add the line:Network Management: solaris:cmd:RO::/usr/sbin/cryptoadm:euid=0

C. To the file /etc/security/exec_attr.d/local-entries, add the line: Network Management:solaris:cmd:::/usr/sbin/cryptoadm:euid=0

D. Assign the sadmin role to alice, for example by running usermod -R sadmin alice.

E. Assign the Network Management profile to the sadmin role, for example by running rolemod -P "Network Management" sadmin.

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: C D E

Explanation:

The required configuration has three links in the Solaris RBAC chain: a command must be added to the rights profile, that rights profile must be assigned to the role, and the user must be authorized to assume that role. An execution-attribute entry belongs in the `exec_attr` database. In Solaris 11, local RBAC entries can be maintained in the corresponding `/etc/security/exec_attr.d/local-entries` file. The entry associates `/usr/sbin/cryptoadm` with the `Network Management` profile and supplies an effective UID of zero, allowing the command to run with the required administrative identity when invoked through the profile.

The `sadmin` role must include the `Network Management` profile, normally configured with a command such as `rolemod -P "Network Management" sadmin`. Alice must also be assigned `sadmin` as an available role, normally with `usermod -R sadmin alice`. After these assignments, Alice can use `su - sadmin` (subject to role authentication) and execute the profile-controlled command. Oracle documents the execution-attribute database and the relationship between profiles, roles, and users in its Solaris RBAC documentation: exec_attr(5) and Oracle Solaris 11.4 documentation.

QUESTION NO: 28

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

Why does the process continue to run?

A. You can kill a process only if you are root.

B. You must own the process or have appropriate privilege to send it a signal.

C. You can kill the process only with the pkill command.

D. You need to kill the process with a stronger kill signal.

ANSWER: B

Explanation:

You must own the process or have appropriate privilege to send it a signal. A process is not terminated merely because a user specifies its process ID and uses signal 9. Before delivering a signal, Solaris checks whether the caller is permitted to signal the target process. Ordinarily, the caller's real or effective user ID must match an eligible user ID of the target process. A user with the applicable process-control privilege can also perform this action. Because User1 is attempting to terminate a process owned by User2, the signal request is denied unless User1 has the required privilege.

Signal 9, `SIGKILL`, is significant because the receiving program cannot catch, block, or ignore it; once Solaris successfully delivers it, the process is forcibly terminated without an opportunity for normal cleanup. However, that behavior does not bypass Solaris authorization checks. The `kill` utility reports the permission failure, and the target process remains running because no signal was delivered. See Oracle's kill(1) documentation and the Solaris privileges(7) documentation for the signal-permission and privilege model.

QUESTION NO: 29

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

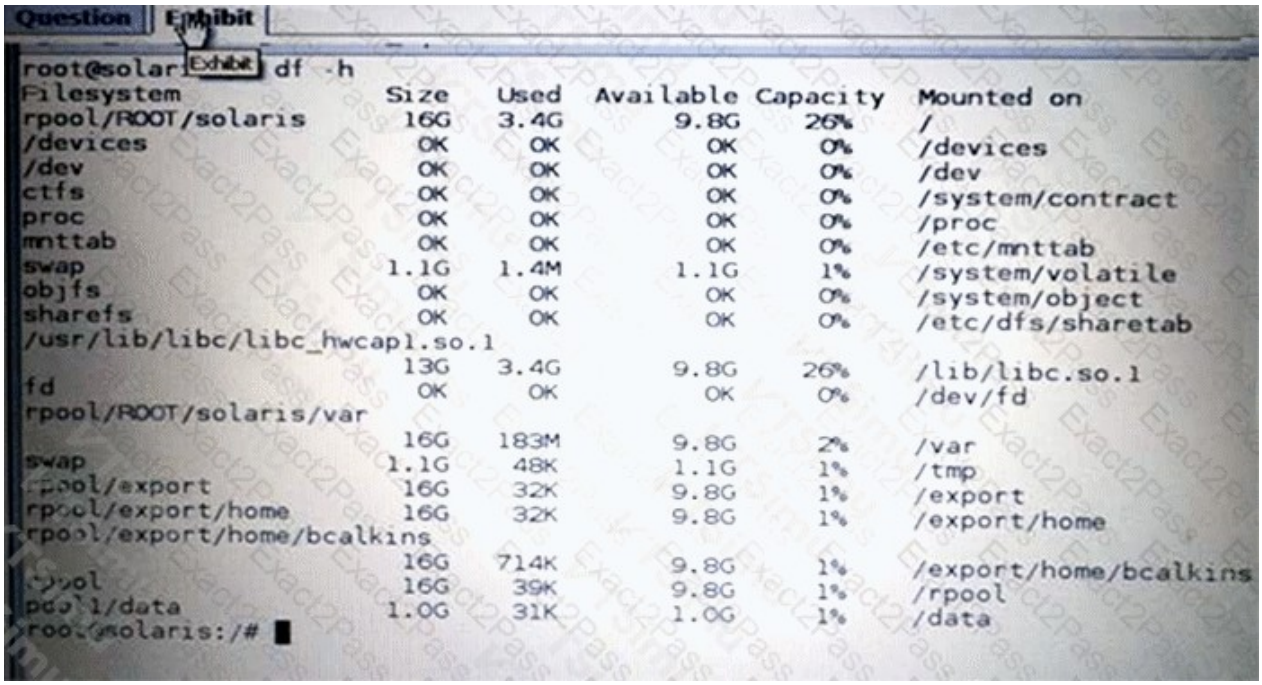

You are configuring a new server. This new server has the following storage pool configured:

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

This new server also has the following file systems configured:

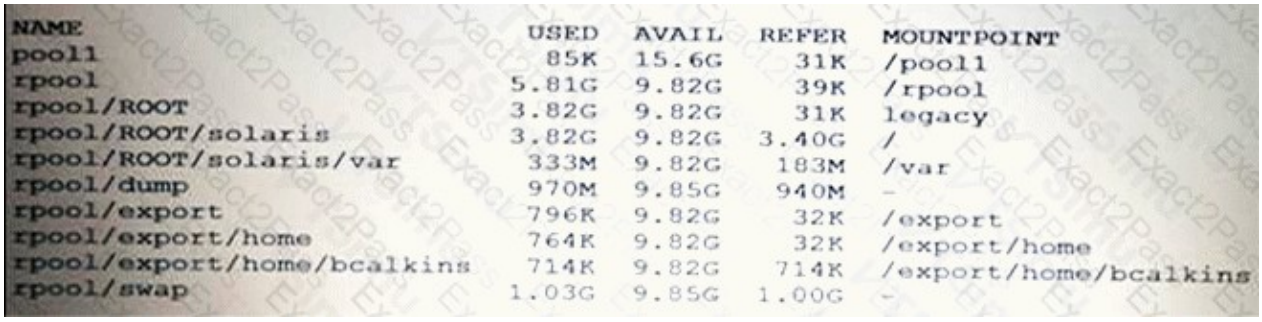

When you are finished building this new server, the pool1/data dataset must be an exact duplicate of note server. What is the correct procedure to create the pool1/data dataset on this new server?

A. zfs create –o mountpoint=/data –o refquota=1g pool1/data

B. zfs set mountpoint=none pool1
zfs create pool1/data

C. zfs set mountpoint=none pool1
zfs create –o mountpoint=/data –o quota=1g pool1/data

D. zfs create quota=1g pool1/data

E. zfs create mountpoint=/data pool1/data

F. zfs set quota=1g pool1/data

ANSWER: C

Explanation:

`zfs set mountpoint=none pool1` followed by `zfs create -o mountpoint=/data -o quota=1g pool1/data` reproduces the relevant ZFS configuration shown for the source system. Setting the parent dataset's `mountpoint` property

to `none` ensures that the pool's root dataset is not automatically mounted. This is important when the required layout is for the child dataset, rather than the pool root, to provide the mounted file system at `/data`. Creating `pool1/data` with `-o mountpoint=/data` explicitly assigns the required mount location at creation time. The `-o quota=1g` property establishes the one-gigabyte ZFS quota for the dataset and its descendants, matching the displayed source dataset property. ZFS supports assigning dataset properties during creation through the `zfs create -o property=value` syntax, and the mountpoint and quota properties are persistent dataset properties. Oracle documents the `zfs create` property syntax and the behavior of ZFS quotas in the zfs(8) manual page and its ZFS file-system administration guide.

QUESTION NO: 30

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:

IP is correct because it is the network-layer protocol that provides logical addressing and enables packets, formally called IP datagrams, to be forwarded between separate networks. Each IP packet contains source and destination IP addresses. Hosts use these addresses to determine whether a destination is local or requires delivery through a router, while routers examine the destination address and forward the packet toward the appropriate next hop. This forwarding behavior is what enables internetworking: communication across multiple distinct physical networks rather than only within one local network segment.

Both IPv4 and IPv6 define the addressing and packet-delivery framework used by routers. IP itself is connectionless and provides best-effort delivery; it does not guarantee that a packet will arrive, arrive only once, or arrive in order. Those characteristics allow IP to serve as the common layer for carrying many kinds of higher-level traffic across network boundaries. Oracle Solaris networking uses IP interfaces, routes, and gateways as the basis for communication with remote networks. For authoritative background, see RFC 791: Internet Protocol and Oracle's Configuring IP Interfaces in Oracle Solaris.

QUESTION NO: 31

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

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

A. svccfg -s svc:/system/system -log setprop start/exec= "syslogd -t"Restart the syslogd daemon.

B. Set the following parameter in the /etc/syslogd.conf file: LOG_FROM_REMOTE= YESRestart the syslogd daemon.

C. svcadm enable svc:/system/system -log/config/log_from_remoteRestart the syslogd daemon.

D. svccfg -s svc:/system/system-log setprop config/log_from_remote=trueRestart the syslogd daemon.

E. Set the following parameter in the /etc/default/syslogd file: LOG_FROM_REMOTE=YESRestart the syslogd daemon.

ANSWER: D E

Explanation:

`svccfg -s svc:/system/system-log setprop config/log_from_remote=true` followed by a restart is correct because the Solaris system-log SMF service exposes `config/log_from_remote` as the service property that controls whether `syslogd` receives log messages from remote hosts. Setting this Boolean property to true enables remote-message reception for the SMF-managed daemon. The service must then be restarted, or refreshed where appropriate, for the running daemon to use the revised configuration.

Setting `LOG_FROM_REMOTE=YES` in `/etc/default/syslogd` and restarting the daemon is also correct. This file contains default settings consumed by `syslogd`; `LOG_FROM_REMOTE` controls remote-message logging. A value of `YES` permits receipt and logging of remote syslog traffic, whereas a disabled setting corresponds to starting the daemon with behavior equivalent to the `-t` option. These settings control acceptance of remote messages; network reachability and any applicable firewall policy must also permit syslog traffic. See the Oracle Solaris syslogd manual page and the Oracle Solaris documentation library.

QUESTION NO: 32

User jack logs in to host solaris and then attempts to log in to host oracle using ssh. He receives the following error message:

The authenticity of host oracle (192.168.1.22) can't be established. RSA key fingerprint is 3B:23:a5:6d:ad:a5:76:83:9c:c3:c4:55:a5:18:98:a6 Are you sure you want to continue connecting (yes/no)?

Which two are true?

A. The public host key supplied by solaria is not known to the host oracle.

B. The error would not occur if oracle's RSA public host key, represented by the fingerprint shown, were added to the `/etc/ssh/ssh_known_hosts` file on solaris.

C. The private host key supplied by oracle is not known to solaris.

D. If jack answers yes, the RSA public key for the host oracle will be added to the known_hosts file for the user jack.

E. The public host key supplied by oracle is not known to the host solaris.

ANSWER: B D

Explanation:

This message is the normal SSH first-contact verification prompt. The SSH client running on `solaris`, on behalf of jack, has received oracle's RSA host public key but does not yet have a trusted saved entry for that host key in the applicable known-hosts databases. If oracle's RSA host public key were pre-recorded in the system-wide `/etc/ssh/ssh_known_hosts` file on solaris, and the entry matched the key presented by oracle, SSH could verify the server without displaying this first-contact authenticity prompt. A known-hosts entry contains the host identity and its public host key; the displayed fingerprint is a compact representation used to verify that key out of band.

If jack responds `yes`, SSH accepts the presented host key and records oracle's RSA public host key in jack's per-user known-hosts file, normally `~/.ssh/known_hosts`. On later connections, SSH compares oracle's presented host key with the saved key. This host-key checking is an important protection against server impersonation and man-in-the-middle attacks, so the fingerprint should be validated through a trusted administrative channel before accepting it. Oracle documents SSH host authentication and known-host files in the Oracle Solaris Secure Shell documentation and the ssh manual page.

QUESTION NO: 33

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

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

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

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

```
ascii name  = <ATA-VBOX HARDDISK-1.0-146.00GB>
bytes/sector =  512
sectors = 306184191
accessible sectors = 306184158
Part      Tag    Flag     First Sector         Size         Last Sector
  0        usr    wm               256      145.99GB           306167774
  1 unassigned    wm                 0           0                   0
  2 unassigned    wm                 0           0                   0
  3 unassigned    wm                 0           0                   0
  4 unassigned    wm                 0           0                   0
  5 unassigned    wm                 0           0                   0
  6 unassigned    wm                 0           0                   0
  8   reserved    wm         306167775      8.00MB           306184158
```

Which two are true regarding the ZFS storage pool?

A. The data on c3t3d0 is duplicated on c3t4d0.

B. The data is striped across disks c3t3d0 and c3t4d0 and mirrored across vdevs mirror-0 and mirror-1.

C. The storage pool is 146 GB total size (rounded to the nearest GB).

D. The storage pool is 584 GB total size (rounded to the nearest GB).

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

ANSWER: A E

Explanation:

The data on c3t3d0 is duplicated on c3t4d0 because those devices form a mirrored top-level virtual device. In a ZFS mirror, each write is stored on every member of that mirror, so either disk contains a complete copy of the data allocated to that vdev. ZFS then combines capacity from separate top-level vdevs by striping allocations across them. Thus, the pool provides both mirrored redundancy within each mirror and aggregate usable capacity across the two mirrored vdevs.

The storage pool is 292 GB total size (rounded to the nearest GB) because each mirror contributes only the capacity of one of its two members. Based on the displayed 146-GB usable partition per disk, the first two-disk mirror contributes approximately 146 GB and the second two-disk mirror contributes approximately 146 GB. The two top-level mirrored vdevs together therefore provide approximately 292 GB of usable pool capacity. ZFS reports and uses capacity at the pool/vdev level; replication in a mirror protects data but does not double usable capacity. Oracle documents mirrored storage-pool configurations and their capacity behavior in the zpool(8) manual page and in the ZFS administration documentation.