Authentic Best resources for 1Z0-821 Test Engine Practice Exam [Q64-Q83

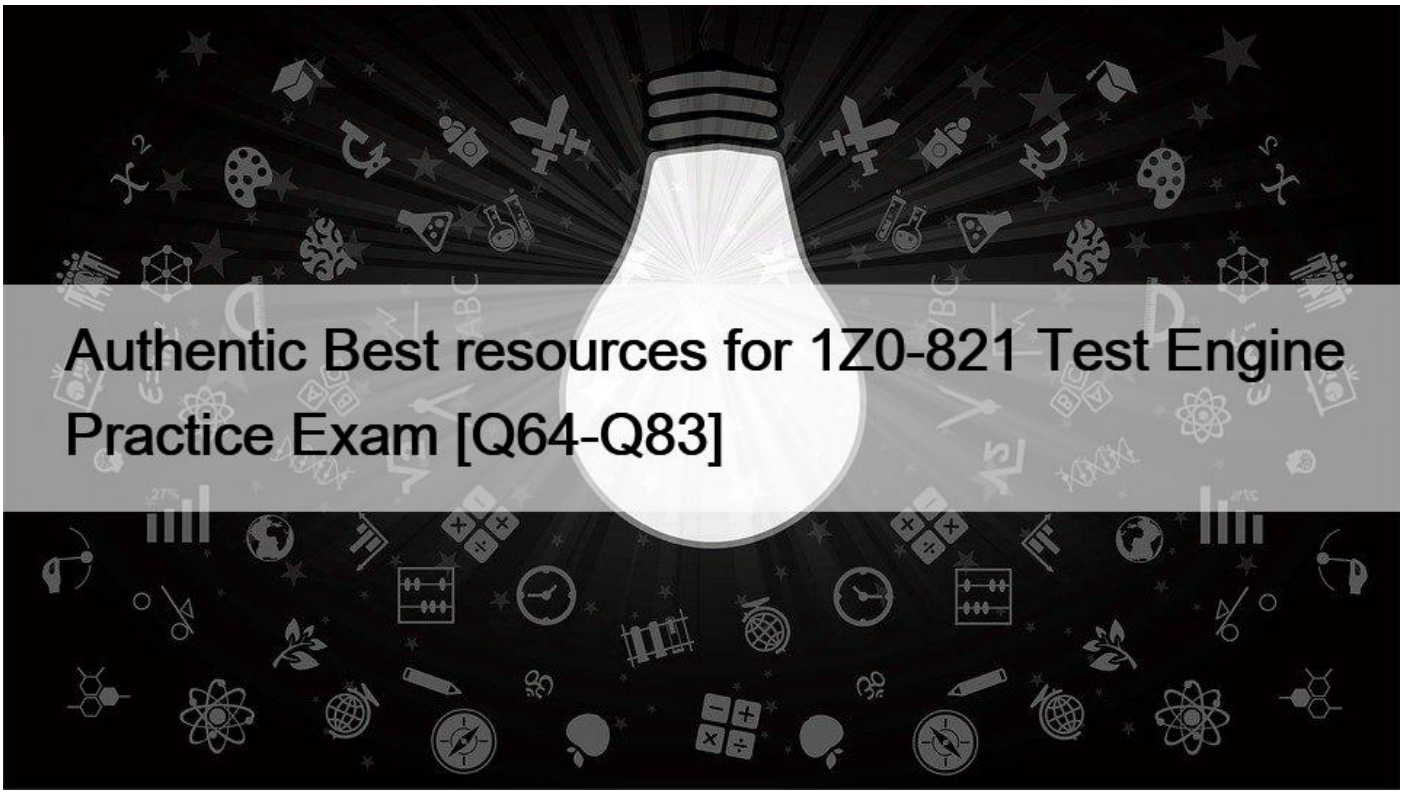

Authentic Best resources for 1Z0-821 Test Engine Practice Exam
[2022] 1Z0-821 PDF Questions - Perfect Prospect To Go With TopExamCollection Practice Exam

QUESTION 64

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.
* Shut the system down, replace disk c3t3d0, and boot the system. When the system is booted, execute the zpool clear pool1 command.

* Shut the system down, replace disk c3t3d0, and boot the system. When the system is booted execute the zpool online pool1 command.
* Shut the system down, replace disk c3t3d0, and boot the system. When the system is booted, execute the zpool replace pool1 c3t3d0 command.
* Shut the system down, replace disk c3t3d0, and boot the system. When the system is booted, execute the zpool replace pool1 c3t3d0 c3t3d0 command.

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

cfgadm -c configure c1::dsk/c1t0d0

zpool replace rpool c1t0d0s0

zpool online rpool c1t0d0s0

zpool status rpoolSPARC# installboot -F zfs /usr/platform/`uname -i`/lib/fs/zfs/bootblk /dev/rdsk/c1t0d0s0 x86# installgrub /boot/grub/stage1 /boot/grub/stage2 /dev/rdsk/c1t9d0s0

QUESTION 65

The core dump configuration in your non global zone is

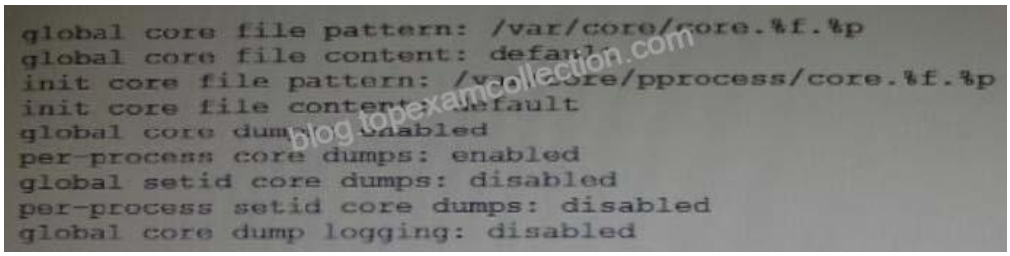

A user is running a process in a non-global zone (testzone) and the process crashes. The process information is: user126632618017:46:42pts/20:00/usr/bin/bash When the user’s process crashes in testzone,a non-global zone,where will the core dump be saved?

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

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.

QUESTION 66

After installing the OS, the following network configuration information is displayed from the system:

Which option describes the state of this server?
* The automatic network configuration option was chosen during the installation of the OS.
* The manual network configuration option was chosen during the installation of the OS.
* The network was not configured during the installation of the OS.
* The network interface is configured with a static IP address.

Only the loopback addresses are configured. No IP address is configured.

QUESTION 67

View the Exhibit.

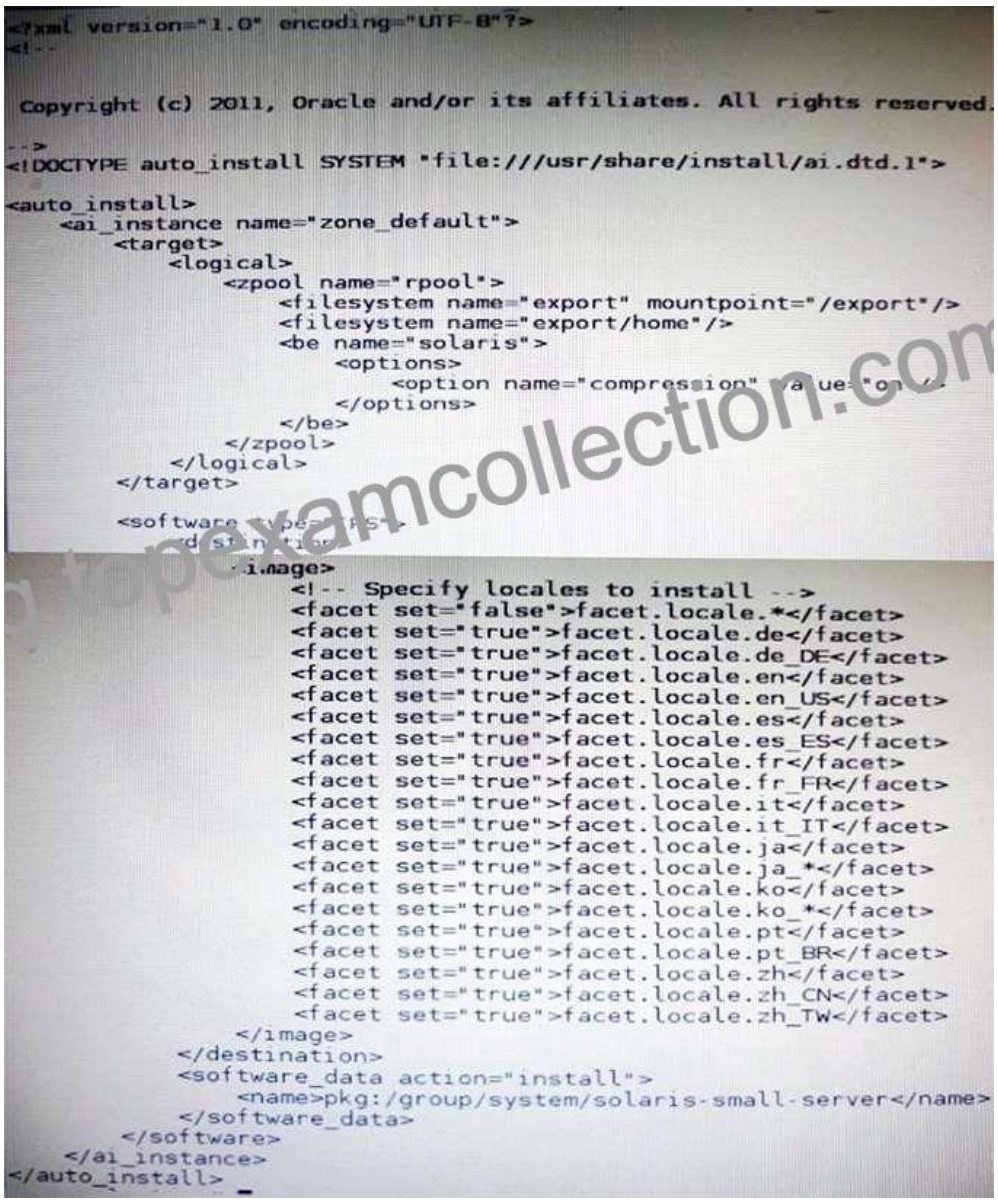

The file came from your Automated Installer (AI) install server.

The file is _____.

* An AI SC profile for non-global zones
* The default AI conf ig file for non-global zones
* The default AI manifest for non-global zones
* A custom AI manifest

ai_manifest

– Automated installation manifest file format

Synopsis

/usr/share/install/ai.dtd.1

Some customizations have been made, such as the selection of specific locales.

QUESTION 68

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

What type of file does this generate?
* a.out
* objdump
* core dump
* tape dump
* crash dump

Explanation/Reference:

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 nonvolitile 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.”

QUESTION 69

You are asked to determine user jack’s default login directory. Which command would provide you with useful information?
* cat /etc/passwd | grep jack
* cat /etc/group | grep jack
* cat /etc/shadow | grep jack
* cat /etc/default/passwd | grep jack

Explanation

Explanation:

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

1. Username: It is used when user logs in. It should be between 1 and 32 characters in length.

2. Password: An x character indicates that encrypted password is stored in /etc/shadow file.

3. User ID (UID): Each user must be assigned a user ID (UID). UID 0 (zero) is reserved for root and UIDs 1-99 are reserved for other predefined accounts. Further UID 100-999 are reserved by system for administrative and system accounts/groups.

4. Group ID (GID): The primary group ID (stored in /etc/group file)

5. User ID Info: The comment field. It allow you to add extra information about the users such as user’s full name, phone number etc. This field use by finger command.

6. Home directory: The absolute path to the directory the user will be in when they log in. If this directory does not exists then users directory becomes /

7. Command/shell: The absolute path of a command or shell (/bin/bash). Typically, this is a shell. Please note that it does not have to be a shell.

QUESTION 70

Choose three options that describe the features associated with a Live Media installation.
* does not allow the root user to log in to the system directly from the console (or any terminal)
* provides a “hands free” installation
* installs the desktop based packages
* can be used to install only x86 platforms
* installs the server-based set of packages only
* allows both automatic and manual configuration of the network
* installs both the server-based and desktop-based package

The graphical installer is officially known as the “Live Media.” This means that Oracle Solaris can be booted into RAM, causing zero impact on your existing operating system.

After it is loaded, you are free to experiment with Oracle Solaris to determine whether it is something you would like to install to your system.

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

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

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

QUESTION 71

The crash dump notification on your server is:

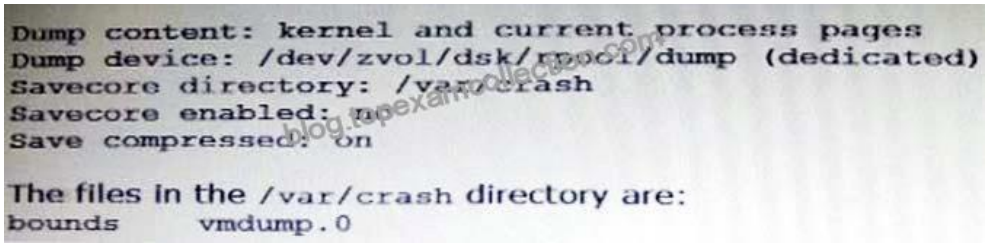

Documentation states that there should be two core files for each crash dump in the /var/crash directory named vmdump.0 Which command should you choose to display theses two files?
* savecore -f vmdump.0
* dumpadm uncompressed
* gunzip vmdump.0
* dumpadm -z off
Explanation/Reference:

Explanation:

Decompress using savecore -f vmdump.0

savecore – save a crash dump of the operating system

-f dumpfile Attempt to save a crash dump from the specified file instead of from the system’s current dump device. This option may be useful if the information stored on the dump device has been copied to an on- disk file by means of the dd(1M) command.

QUESTION 72

Which command should you choose to display the current parameters for the FSS scheduler?
* dispadmin – c FSS
* prionctl -c FSS
* dispadmin -c FSS -g
* priocntl -c FSS -g
Explanation/Reference:

Explanation:

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

-c class

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

-g

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

QUESTION 73

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

Which command would display the Solaris 11 build version currently running on your system?
* pkg info all
* cat /etc/release
* cat /etc/update

* prtconf | grep -i update
* pkg info entire

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

cat /etc/release

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

Example: # cat /etc/release Oracle Solaris 10 8/11 s10s_u10wos_17b SPARC Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved. Assembled 23 August 2011

QUESTION 74

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

* A file on a UFS file system
* A Veritas Volume Manager (VxVM) volume
* A LUN In a hardware RAID array
* A disk slice from an SMI labeled disk
* A Solaris Volume Manager (SVM) volume
* An EFI labeled disk

Explanation/Reference:

Explanation:

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

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

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

QUESTION 75

You need to connect two nonglobal zones using a private virtual network.

Identify the network resources required in the global zone to accomplish this.

* an etherstub and two virtual network interfaces
* a virtual bridge
* two virtual network interfaces.
* two etherstubs

QUESTION 76

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

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

QUESTION 77

Which command would you use from the bash shell to determine the total amount of physical memory installed in your Solaris system (x86 and SPARC)?
* uname -a
* prtconf | grep -i memory
* sysdef | grep -i memory
* vmstat
* prtdiag | grep -i memory

The prtconf command prints the system configuration informa-tion. The output includes the total amount of memory,and the configuration of system peripherals formatted as a dev-ice tree.

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

QUESTION 78

You wish to edit your crontab file that is located in /var/spool/cron/crontab. What command must you enter to edit this file?
* crontab-e
* crontab-e/var/spool/cron/crontab
* crontab-r
* crontab-e/etc/default/cron

The main tool for setting up cron jobs is the crontab command, though this is not available on every Unix variant. Typically under Solaris or Linux one would create a new crontab or edit an existing one, using the command; crontab -e

Use the ls -l command to verify the contents of the/var/spool/cron/crontabs file.

Reference: System Administration Guide: Advanced Administration, Creating and Editing crontab Files

QUESTION 79

View the Exhibit to inspect the boot environment Information displayed within a non global zone on your system.

Which two options describe the solaris-1 boot environment?

* The solaris-1 boot environment is not bootable.
* The solaris-1 boot environment is incomplete.
* The solaris-1 boot environment was created automatically when the non global zone was created.
* The solaris-1 boot environment was created in the non-global zone using the beadm create command.
* The solaris-1 boot environment is associated with a non active global zone boot environment.

Explanation/Reference:

Explanation:

A: The – of the Active Column indicates that this boot environment is inactive, and hence not bootable.

Note: The values for the Active column are as follows:

R – Active on reboot.

N – Active now.

NR – Active now and active on reboot.

“-” – Inactive.

“!” – Unbootable boot environments in a non-global zone are represented by an exclamation point.

D: beadm create

Creates a new boot environment name, beName.

Note: beadm list

Lists information about the existing boot environment, which is be Name, or lists information for all boot environments if be Name is not provided.

Note: Using beadm Utility (Tasks)

You can use the beadm utility to create and manage snapshots and clones of your boot environments.

Note the following distinctions relevant to boot environment administration:

* A snapshot is a read-only image of a dataset or boot environment at a given point in time. A snapshot is not bootable.

* A boot environment is a bootable Oracle Solaris environment, consisting of a root dataset and, optionally, other datasets mounted underneath it. Exactly one boot environment can be active at a time.

* A clone of a boot environment is created by copying another boot environment. A clone is bootable.

QUESTION 80

User brian changes the permissions for db_data this command:

chmod 4755 db_data

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

Explanation/Reference:

Explanation:

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

Here we do not know if brian owns db_data.

Note:

Permission

7 full

6 read and write

5 read and execute

4 read only

3 write and execute

2 write only

1 execute only

0 none

0 — no permission

1 –x execute

2 -w- write

3 -wx write and execute

4 r– read

5 r-x read and execute

6 rw- read and write

7 rwx read, write and execut

Solaris: Solaris Advanced User’s Guide

QUESTION 81

User jack on host solaris attempts to use ssh to log in to host oracle and receives this message:

jack@solaris:~$ ssh oracle

ssh: connect to host oracle port 22: connection refused

What is the problem?
* Host oracle does not have a valid host public key.
* Host oracle does not have a valid host private key.
* Host solaris does not have a valid host public key.
* Host does not have a valid host private key.
* Host solaris is not configured for host-based authentication.
* Host oracle is not configured for host-based authentication.
* Host oracle is not running the ssh service.
* Host solaris is not running the ssh service.

Explanation/Reference:

Explanation:

The host he is trying to connect to (oracle) is not running the required service (ssh).

QUESTION 82

When upgrading an existing system from Solaris 11 Express to Oracle Solaris 11, what happens to the datalink names?
* They follow the default naming convention for the newly installed version.
* They maintain their names.
* They are called eth#.
* They are called el00g#.
* They are left unnamed, to avoid conflicts, and need to be renamed after the installation process is complete.

Explanation/Reference:

Explanation:

Network configuration in Oracle Solaris 11 includes

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

Note: There is no upgrade path from Oracle Solaris 10 to Oracle Solaris 11. You must perform a fresh installation.

QUESTION 83

Your mentor suggests using the dladm rename-link command to rename the network datalinks.

What are the two advantages of following this advice?
* It can clarify which network interface has what purpose.
* It can simplify specifying the network interface with the dladm modify-aggr command.
* It can simplify specifying the network interface with the dladm modify-bridge command.
* It can simplify IP filter rule changes if the network interface is replaced with a different type.
* It can prevent accidental deletion of the network interface with the dladm delete-phys command.
* It can prevent accidental deletion of the network interface configuration with the ipadm delete-addr command.

A: To rename the bge0 link to mgmt0,enter the following command:

dladm rename-link bge0 mgmt0

E: Consider that the bge0 device,whose link was named mgmt0 as shown in the previous example,needs to be replaced with a ce0 device because of a hardware failure. The bge0 NIC is physically removed,and replaced with a new ce0 NIC. To associate the newly added ce0 device with the mgmt0 configuration previously associated with bge0,enter the following command:

dladm rename-link ce0 mgmt0

Note: How to Rename a Datalink

Use this procedure if you want to change a datalink name to a customized name. For example,some of the datalinks in upgraded system might have retained legacy hardware-based names and you want to change these names to generic ones.

Note: dladm rename-link [-R root-dir] link new-link Rename link to new-link. This is used to give a link a meaningful name,or to

associate existing link configuration such as link properties of a removed device with a new device.