Oracle

Exam Questions 1Z0-821

Oracle Solaris 11 System Administrator

NEW QUESTION 1
To assist in examining and debugging running processes, Solaris 11 has a utility that returns pro arguments and the names and values of environment variables.
What is the name of this utility?

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

Answer: B

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

NEW QUESTION 2
To confirm the IP addresses and netmasks have been correctly configured on the network interfaces, which command(s) should you use?

A. ipadm show-if
B. ipadm show-nic
C. ipadm show-addr
D. ipadm show-addripadm show-mask
E. ipadm show-ipipadm show-mask
F. ipadm show-config

Answer: C

Explanation:
Show address information, either for the given addrobj or all the address objects configured on the specified interface, including the address objects that are only in the persistent configuration.
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 3
You have been asked to do an orderly shutdown on a process with a PID of 1234, with the kill command.
Which command is best?

A. kill -2 1234
B. kill -15 1234
C. kill -9 1234
D. kill -1 1234

Answer: B

Explanation:
On POSIX-compliant platforms, SIGTERM is the signal sent to a process to request its termination. The symbolic constant for SIGTERM is defined in the header file signal.h. Symbolic signal names are used because signal numbers can vary across platforms, however on the vast majority of systems, SIGTERM is signal #15.
SIGTERM is the default signal sent to a process by the kill or killall commands. It causes the termination of a process, but unlike the SIGKILL signal, it can be caught and interpreted (or ignored) by the process. Therefore, SIGTERM is akin to asking a process to terminate nicely, allowing cleanup and closure of files. For this reason, on many Unix systems during shutdown, init issues SIGTERM to all processes that are not essential to powering off, waits a few seconds, and then issues SIGKILL to forcibly terminate any such processes that remain.

NEW QUESTION 4
Which command would you use from the bash shell to determine the total amount of physical memory installed in your Solaris system (x86 and SPARC)?

A. uname –a
B. prtconf | grep –i memory
C. sysdef | grep –i memory
D. vmstat
E. prtdiag | grep –i memory

Answer: B

Explanation:
The prtconf command prints the system configuration information. The output includes the total amount of memory, and the configuration of system peripherals formatted as a device tree.
If a device path is specified on the command line for those command options that can take a device path, prtconf will only display information for that device node.

NEW QUESTION 5
You need to update an OS image on a client. The pkg publishers command displays the wrong publisher with the wrong update:
PUBLISHERTYPESTATUSURI
Solaris origin onlinehttp://pkg.oracle.com/solaris/release
The update is available on the updated publisher: PUBLISHERTYPESTATUSURI
Solaris originonlinehttp://sysA.example.com
Select the option that describes the procedure used to update the OS image on the system from the updated publisher.

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

Answer: C

Explanation:
You can use the pkg set-publisher command to change a publisher URI. Changing a Publisher Origin URI
To change the origin URI for a publisher, add the new URI and remove the old URI. Use the -g option to add a new origin URI. Use the -G option to remove the old origin URI.
pkg set-publisher -g http://pkg.example.com/support \
-G http://pkg.example.com/release example.com
Note: You can use either the install or update subcommand to update a package.
The install subcommand installs the package if the package is not already installed in the image. If you want to be sure to update only packages that are already installed, and not install any new packages, then use the update subcommand.

NEW QUESTION 6
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 <zonepath>/root directory before booting the non- global zone.

Answer: C

Explanation:
Oracle Solaris 10 branded zones – Oracle Solaris 10 Zones provide an Oracle Solaris 10
environment on Oracle Solaris 11. You can migrate an Oracle Solaris 10 system or zone to a solaris10 zone on an Oracle Solaris 11 system in the following ways:
* Create a zone archive and use the archive to create an s10zone on the Oracle Solaris 11 system.
This option applies in the current scenario.
Example of command to Install the Oracle Solaris 10 non-global zone. s11sysB# zoneadm -z s10zone install -u -a /pond/s10archive/s10.flar
* Detach the zone from the Oracle Solaris 10 system and attach the zone on the Oracle Solaris 11 zone. The zone is halted and detached from its current host.
The zonepath is moved to the target host, where it is attached.
Note:
install [-x nodataset] [brand-specific options] A subcommand of the zoneadm.
Install the specified zone on the system. This subcommand automatically attempts to verify first. It refuses to install if the verify step fails.
-u uuid-match
Unique identifier for a zone, as assigned by libuuid(3LIB). If this option is present and the argument is a non-empty string, then the zone matching the UUID is selected instead of the one named by the -z option, if such a zone is present.

NEW QUESTION 7
Which two options are accurate regarding the non-global zone console?

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

Answer: AC

Explanation:
A: How to Log In to the Zone Console Use the zlogin command with the -C option and the name of the zone, for example, my-zone.
global# zlogin -C my-zone
C: To disconnect from a non-global zone, use one of the following methods.
* To exit the zone non-virtual console: zonename# exit
* To disconnect from a zone virtual console, use the tilde (~) character and a period: zonename# ~.

NEW QUESTION 8
User jack makes use of the bash shell; his home directory is/export/home/jack.
What is the correct setting of umask, and where should it be set, to allow jack to create a shell script using the vi editor, that is executable by default?

A. It is not possible to make a script executable without using the chmod command.
B. umask value of 0002 set in /etc/profile
C. umask value of 0002 set in /export/home/jack/.bashrc
D. umask value of 0722 set in /etc/profile

E. umask value of 0722 set In /export/home/jack/.bashrc

Answer: B

Explanation:
The user file-creation mode mask (umask) is use to determine the file permission for newly created files. It can be used to control the default file permission for new files. It is a four- digit octal number.
You can setup umask in /etc/bashrc or /etc/profile file for all users. By default most Unix distro set it to 0022 (022) or 0002 (002).
1. The default umask 002 used for normal user. With this mask default directory permissions are 775 and default file permissions are 664.
2. The default umask for the root user is 022 result into default directory permissions are 755 and default file permissions are 644.
3. For directories, the base permissions are (rwxrwxrwx) 0777 and for files they are 0666 (rw-rw-rw).
In short,
1. A umask of 022 allows only you to write data, but anyone can read data.
2. A umask of 077 is good for a completely private system. No other user can read or write your data if umask is set to 077.
3. A umask of 002 is good when you share data with other users in the same group. Members of your group can create and modify data files; those outside your group can read data file, but cannot modify it. Set your umask to 007 to completely exclude users who are not group members.

NEW QUESTION 9
You are troubleshooting the failure of a computer to mount an NFS file system hosted by a server (hostname mars) in the local area network.
Select the three commands that will enable you to identify the problem.

A. ping - s mars
B. cat /etc/vfstab
C. cat /etc/dfs/dfstab
D. sharemgr show -v
E. showmount -e mars
F. rpcinfo -s mars | egrep ‘nfs|mountd’

Answer: BEF

Explanation:
B: The mount point Error. The following message appears during the boot process or in response toan explicit mount request and indicates a non-existent mount point.
Mount: mount-point /DS9 does not exist.
To solve the mount point error condition, check that the mount point exists on the client. Check the spelling of the mount point on the command line or in the /etc/vfstab file (B) on the client, or comment outthe entry and reboot the system.
Note: The /etc/vfstab file lists all the file systems to be automatically mounted at system boot time, with the exception of the /etc/mnttab and /var/run file systems.
E: showmount
This command displays all clients that have remotely mounted file systems that are shared from an NFS server, or only the file systems that are mounted by clients, or the shared file systems with the client access information. The command syntax is:
showmount [-ade] [hostname]
where -a prints a list of all the remote mounts (each entry includes the client name and the
directory), -d prints a list of the directories that are remotely mounted by clients, -e prints a list of the files shared (or exported), and hostname selects the NFS server to gather the information from. If hostname is not specified the local host is queried.
F: * mountd Daemon
This daemon handles file-system mount requests from remote systems and provides access control. The mountd daemon checks /etc/dfs/sharetab to determine which file systems are available for remote mounting and which systems are allowed to do the remote mounting.
* Commands for Troubleshooting NFS Problems
These commands can be useful when troubleshooting NFS problems. rpcinfo Command
This command generates information about the RPC service that is running on a system.

NEW QUESTION 10
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-addripadm show-mask

Answer: C

Explanation:
Show address information, either for the given addrobj or all the address objects configured on the specified interface, including the address objects that are only in the persistent configuration.
State can be: disabled, down, duplicate, inaccessible, ok, tentative Example:
ipadm show-addr
ADDROBJ TYPE STATE ADDR
lo0/v4 static ok 127.0.0.1/8 lo0/v6 static ok ::1/128

NEW QUESTION 10
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 13
You are logged in as root to a newly installed Solaris 11 system. You issue the command useradd -d, and then examine the /usr/sadm/defadduser file. This file includes the entry defshell=/bin/sh. Which shell will now be the default for the next account created?

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

Answer: A

Explanation:
Oracle Solaris 11 introduces user environment and command-line argument changes that include the following:
* Shell changes - The default shell, /bin/sh, is now linked to ksh93. The default user shell is the Bourne-again (bash) shell.
* The legacy Bourne shell is available as /usr/sunos/bin/sh.
* The legacy ksh88 is available as /usr/sunos/bin/ksh from the shell/ksh88 package.
* Korn shell compatibility information is available in /usr/share/doc/ksh/COMPATIBILITY.

NEW QUESTION 15
User jack, whose account is configured to use the korn shell, logs in and examines the value of his PATH environment variable:

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

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

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

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

```
jack@solaris: useradd
```

What will happen, and why?

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

Answer: D

NEW QUESTION 16
Which two are user definable OpenBoot parameters that can be set in the OpenBoot PROM?

A. IP address for the system console
B. Host ID
C. System date and time
D. Default boot device
E. Verbose hardware diagnostics
F. Powering off the hardware

Answer: DE

Explanation:
The NVRAM chip stores user-definable system parameters, also referred to as NVRAM variables or EEPROM parameters. The parameters allow administrators to control
variables such as the default boot device and boot command. The NVRAM also contains writeable areas for user-controlled diagnostics, macros, and device aliases. NVRAM is where the system identification information is stored, such as the host ID, Ethernet address, and time-of-day (TOD) clock.
Examples of NVRAM variables:
Variable Default Description boot-device disk or net The device from which to start up.
diag-device net The diagnostic startup source device.
diag-file Empty string Arguments passed to the startup program in diagnostic mode. diag-switch? false Whether to run in diagnostic mode

NEW QUESTION 17
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 promp
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 18
You want to delete the IPv4 address on the interface net3. Which command should you use?

A. ipadm delete-ip net3/v4
B. ipadm down-addr net3/v4
C. ipadm disable-if net3/v4
D. ipadm delete-vni net3/v4
E. ipadm delete-addr net3/v4
F. ipadm deiete-ipv4 ner3/v4

Answer: E

Explanation:
The ipadm delete-addr subcommand removes addresses from interfaces. To remove an address from the IPMP group, type the following command:
ipadm delete-addr addrobj
The addrobj uses the naming convention inder-interface/user-string.

NEW QUESTION 22
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 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.”

NEW QUESTION 23
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 25
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 27
The following image properties are displayed on your system:

PROPERTY	VALUE
be-policy	always-new
ca-path	/etc/openssl/certs
check-certificate-revocation	False
flush-content-cache-on-success	True
mirror-discovery	False
preferred-authority	
publisher-search-order	[‘solaris’]
send-uuid	True
signature-policy	verify
signature-required-name	[]
trust-anchor-directory	etc/certs/CA
use-system-repo	False

Which two options describe the boot environment policy property that is currently set for this image?

A. All package operations are performed in a new BE set as active on the next boot.
B. Do not create a new B
C. The install, update, uninstall, or revert operation is not performed if a new BE is required.
D. If a BE is created, do not set it as the active BE on the next boot
E. A reboot is required for all package operations
F. A reboot is not required after a package operation.
G. For package operations that require a reboot, this policy creates a new BE set as active on the next boot.

Answer: DF

Explanation:
Image properties described below.
* be-policy
Specifies when a boot environment is created during packaging operations. The following values are allowed:
/ default
Apply the default BE creation policy: create-backup.
/ always-new (D, F)
Require a reboot for all package operations (D) by performing them in a new BE set as active on the next boot (F). A backup BE is not created unless explicitly requested.
This policy is the safest, but is more strict than most sites need since no packages can be added without a reboot.

NEW QUESTION 29
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 34
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 35
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 39
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 = 256-ccm - o keysource = raw, file : ///my key 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:
Example: Encrypting a ZFS File System by Using a Raw Key
In the following example, an aes-256-ccm encryption key is generated by using the pktool command and is written to a file, /cindykey.file.
pktool genkey keystore=file outkey=/cindykey.file keytype=aes keylen=256
Then, the /cindykey.file is specified when the tank/home/cindy file system is created.
zfs create -o encryption=aes-256-ccm -o keysource=raw, file:///cindykey.file tank/home/cindys

NEW QUESTION 41
Which two accurately identify features of a Solaris 10 branded zone?

A. executes in a Solaris 10 global zone
B. is created by importing a Solaris 10 flash archive
C. enables Linux binary applications to run unmodified
D. provides a complete runtime environment for Solaris 9 applications
E. allows a Solaris 10 global zone to be migrated into a Solaris 10 non-global zone on a Solaris 11 system

Answer: BE

Explanation:
B: It can be created by importing a Solaris 10 flash archive.
You can use the Oracle Solaris Flash archiving tools to create an image of an installed
system that can be migrated into a zone.
The system can be fully configured with all of the software that will be run in the zone before the image is created. This image is then used by the installer when the zone is installed.
Note: You can use alternate methods for creating the archive. The installer can accept the following archive formats:
* cpio archives
* gzip compressed cpio archives
* bzip2 compressed cpio archives
* pax archives created with the -x xustar (XUSTAR) format Ä ufsdump level zero (full) backups
Note:
Branded zones that run an environment different that the OS release on the system
* The lx branded zone introduced in the Solaris 10 8/07 release provides a Linux environment for your applications and runs on x86 and x64 machines on the Oracle Solaris 10 OS.
* The solaris8 and solaris9 branded zones enable you to migrate an Oracle Solaris 8 or Oracle Solaris 9 system to an Oracle Solaris 8 or Oracle Solaris 9 Container on a host running the Oracle Solaris 10 8/07 Operating System or later Oracle Solaris 10 release.
* The Oracle Solaris 10 Container brand is available in OpenSolaris build 127. These branded zones host Oracle Solaris 10 user environments.
Note: One of the powerful features of Solaris 11 is the ability to run a Solaris 10 environment in a zone. Solaris 10 allows you to run Solaris 8 and 9 environments in zones, but only on SPARC.

NEW QUESTION 46
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 p7sip
C. pkg rebuild-index p7zip
D. pkg revert p7zip
E. pkg uninstsll p7zip
F. pkg install p7zip
G. pkg install --no-backup-be p7sip
H. pkg refresh p7zip

Answer: D

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

NEW QUESTION 50
You are the administrator of a system that a large number of developers work on. These developers crash the system, and their applications, on a regular basis. What command would you use to configure where the core files are saved?

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

Answer: E

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

NEW QUESTION 54
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:
No OpenBoot Environment on the Intel Platform. The Intel environment has no OpenBoot PROM or NVRAM. On Intel systems, before the kernel is started, the system is controlled by the basic input/output system (BIOS), the firmware interface on a PC. Therefore, many features provided by OpenBoot are not available on Intel systems.
Note: The Open Boot PROM (OBP) bootloader only exists within SPARC. Before Solaris 10 01/06, the bootloader for Solaris x86 was a Sun customized bootstrap software. After Solaris 10 01/06, it uses GRUB, a well known bootloader that's commonly used in the
Linux world.
With GRUB, it's much easier to make the system dual-boot Linux and Solaris. GRUB extends the capabilities of the bootloader that was not available previously such as the ability to boot from a USB DVD drive. Those who have used Linux will be quite familiar with GRUB and its options.

NEW QUESTION 58
View the Exhibit.

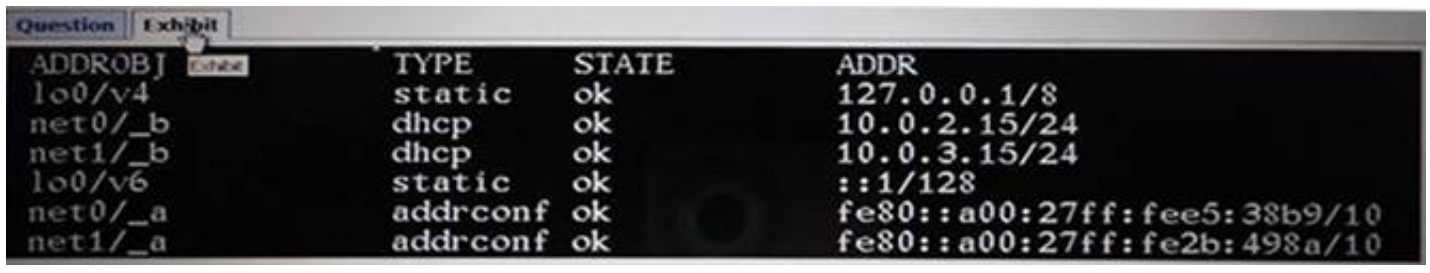

After Installing the OS, you need to verify the network interface information. Which command was used to display the network interface information in the exhibit?

A. ifconfiq –a
B. ipadm show-addr
C. svcs –1 network/physical
D. netstat –a

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 61
You run the command dlstat show-link -r.

Select the two correct statements regarding the information displayed in the INTRS column.

A. No value is listed for virtual network interfaces.
B. A value of 0 is listed for virtual interfaces and ether stubs.
C. The number of Interrupts is listed, which indicates network efficiency.
D. A number equal to the number of transmitted Ethernet frames is listed for physical links.
E. The number of packets that were interrupted by a collision is listed, which may indicate hardware problems.

Answer: CE

Explanation:
In this output, the statistics for interrupt (INTRS) are significant. Low interrupt numbers indicate greater efficiency in performance. If the interrupt numbers are high, then you might need to add more resources to the specific link.
Example:
dlstat -r -i 1
LINK IPKTS RBYTES INTRS POLLS CH<10 CH10-50 CH>50 e1000g0 101.91K 32.86M 87.56K 14.35K 3.70K 205 5
nxge1 9.61M 14.47G 5.79M 3.82M 379.98K 85.66K 1.64K vnic1 8 336 0 0 0 0 0
e1000g0 0 0 0 0 0 0 0
nxge1 82.13K 123.69M 50.00K 32.13K 3.17K 724 24
vnic1 0 0 0 0 0 0 0
Note: dlstat show-link [-r [-F] | -t] [-i interval] [-a] [-p] [-o field[, ...]] [-u R|K|M|G|T|P] [link] Display statistics for a link.
-r
Display receive-side statistics only. Includes bytes and packets received, hardware and software drops, and so forth.
List of supported RX fields: link
iusedby
ibytes ipkts intrs polls
hdrops: hardware drops
sdrops: software drops (owing to bandwidth enforcement) ch<10: number of packet chains of length < 10
ch10-50: number of packet chains of length between 10 and 50 ch>50: number of packet chains of length > 50

NEW QUESTION 62
When upgrading an existing system from Solaris 11 Express to Oracle Solaris 11, what happens to the datalink names?

A. They follow the default naming convention for the newly installed version.
B. They maintain their names.
C. They are called eth#.
D. They are called el00g#.
E. They are left unnamed, to avoid conflicts, and need to be renamed after the installation process is complete.

Answer: A

Explanation:
Network configuration in Oracle Solaris 11 includes
* Generic datalink name assignment – Generic names are automatically assigned to datalinks using the net0, net1, netN naming convention, depending on the total number of network devices that are on the system
Note: There is no upgrade path from Oracle Solaris 10 to Oracle Solaris 11. You must perform a fresh installation.

NEW QUESTION 66
Review the boot environment information displayed on your system:

```
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 two options accurately describe the newBE boot environment?

A. It cannot be destroyed.
B. It cannot be activated.
C. It cannot be renamed.
D. You can create a snapshot of it.
E. It is activated but unbootable.
F. It has been deleted and will be removed at the next reboot.

Answer: BC

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

NEW QUESTION 67
You need to make sure that all of the software packages on your server are up to date. Without installing any updates, which two commands would display .my software updates that are available in the default Oracle repository?

A. pkg list -u

B. pkg verify –u '*'
C. pkg search –u
D. pkg info –r '*'
E. pkg install –nv
F. pkg update –nv '*'

Answer: AD

Explanation:
A: the pgk list command display a list of packages in the current image, including state and other information. By default, package variants for a different architecture or zone type are excluded.
D: pkginfo displays information about software packages that are installed on the system (with the first synopsis, with -l) or that reside on a particular device or directory (with the second synopsis, with -r).
Without options, pkginfo lists the primary category, package instance, and the names of all completely installed and partially installed packages. It displays one line for each package selected.
With -r, retrieve the data from the repositories of the image's configured publishers. Note that you must specify one or more package patterns in this case.

NEW QUESTION 71
A local repository is available on this system and you need to enable clients to access this repository via HTTP. The repository information is:
PUBLISHERTYPESTATUSURI
solarisoriginonlinehttp://sysA.example.com
Identify two of the steps that are required to make the local repository on this server available to the client via HTTP.

A. On the server: set the pkg/inst_root and pkg/readonly properties for the svc:/application/pkg/server:default service and enabled the service
B. On the server: set the sharefs property on the ZFS file system containing the IPS repository.
C. On the client: reset the origin for the solaris publisher.
D. On the client: set the pkg/inst_root and pkg/readonly properties for the svc:/application/server:default service enable the service.
E. On the client: start the pkg.depotd process.

Answer: AE

Explanation:
A: Configure the Repository Server Service
To enable clients to access the local repository via HTTP, enable the application/pkg/server Service Management Facility (SMF) service.
svccfg -s application/pkg/server setprop pkg/inst_root=/export/repoSolaris11
svccfg -s application/pkg/server setprop pkg/readonly=true
E: Use pkg.depotd to serve the repository to clients. Start the Repository Service
Restart the pkg.depotd repository service.
svcadm refresh application/pkg/server
svcadm enable application/pkg/server
To check whether the repository server is working, open a browser window on the localhost location.

NEW QUESTION 76
The advantage of core tiles 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 77
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:
IPv6 defines three address types: unicast
Identifies an interface of an individual node.
multicast
Identifies a group of interfaces, usually on different nodes. Packets that are sent to the multicast address go to all members of the multicast group.
anycast

Identifies a group of interfaces, usually on different nodes. Packets that are sent to the anycast address go to the anycast group member node that is physically closest to the sender.

NEW QUESTION 79
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 83
View the Exhibit.

```
<?xml version="1.0" encoding="UTF-8"?>
<!--

 Copyright (c) 2011, Oracle and/or its affiliates. All rights reserved.

-->
<!DOCTYPE auto_install SYSTEM "file:///usr/share/install/ai.dtd.1">

<auto_install>
    <ai_instance name="zone_default">
        <target>
            <logical>
                <zpool name="rpool">
                    <filesystem name="export" mountpoint="/export"/>
                    <filesystem name="export/home"/>
                    <be name="solaris">
                        <options>
                            <option name="compression" value="on"/>
                        </options>
                    </be>
                </zpool>
            </logical>
        </target>

        <software type="IPS">
            <destination>
                <image>
                    <!-- Specify locales to install -->
                    <facet set="false">facet.locale.*</facet>
                    <facet set="true">facet.locale.de</facet>
                    <facet set="true">facet.locale.de_DE</facet>
                    <facet set="true">facet.locale.en</facet>
                    <facet set="true">facet.locale.en_US</facet>
                    <facet set="true">facet.locale.es</facet>
                    <facet set="true">facet.locale.es_ES</facet>
                    <facet set="true">facet.locale.fr</facet>
                    <facet set="true">facet.locale.fr_FR</facet>
                    <facet set="true">facet.locale.it</facet>
                    <facet set="true">facet.locale.it_IT</facet>
                    <facet set="true">facet.locale.ja</facet>
                    <facet set="true">facet.locale.ja_*</facet>
                    <facet set="true">facet.locale.ko</facet>
                    <facet set="true">facet.locale.ko_*</facet>
                    <facet set="true">facet.locale.pt</facet>
                    <facet set="true">facet.locale.pt_BR</facet>
                    <facet set="true">facet.locale.zh</facet>
                    <facet set="true">facet.locale.zh_CN</facet>
                    <facet set="true">facet.locale.zh_TW</facet>
                </image>
            </destination>
            <software_data action="install">
                <name>pkg:/group/system/solaris-small-server</name>
            </software_data>
        </software>
    </ai_instance>
</auto_install>
```

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

A. An AI SC profile for non-global zones
B. The default AI conf ig file for non-global zones
C. The default AI manifest for non-global zones
D. A custom AI manifest

Answer: D

Explanation:
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.

NEW QUESTION 86
View the Exhibit.

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

format>
```

Which is true regarding the disk drive?

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

Answer: A

Explanation:
Installing a ZFS Root Pool
The installer searches for a disk based on a recommended size of approximately 13 GB.

NEW QUESTION 87
Review the non-global zone configuration displayed below:

```
zonename:  dbzone
zonepath:     /export/dbzone
brand:    Solaris
autoboot:    false
bootargs:
file-mac-profile:
pool:
limitpriv:
scheduling-class:
ip-type: exclusive
hostid:
fs-allowed:
anet:
                linkname: net0
                lower-link: auto
                allowed-address not specified
                configure-allowed-address: true
                defrouter not specified
                allowed-dhcp-cids not specified
                link-protection: mac-nospoof
                mac-address: random
                auto-mac-address: 2:8:20:97:40:20
                mac-preffix not specified
                mac-slot not specified
                vlan-id not specified
                priority nor specified
                rxrings not specified
                rxrings not specified
                mtu not specified
                maxlow not specified
                rxfanout not specified
```

The global zone has 1024 MB of physical memory. You need to limit the non-global zone so that it uses no more than 500 MB of the global zone's physical memory. Which option would you choose?

```
A.
From the global zone, issue these commands:
zonecfg  -z   dbzone
       zonecfq:dbzone>  set  zone.max  -memory=500m

B.
From the global zone, issue these commands:
zonecfq -z  dbzone
     zonecfg:dbzone>add  rctl
     zonecfg:dbzone>  set  zone.max -memory=500m
     zonecfg:dbzone:capped-memory>  end

C.
From the global zone, issue these commands:
zonecfg -z  dbzone
     zonecfg:dbzone> add capped-memory
     zonecfg:dbzone:capped-memory> set physical=500m
     zonecfg:dbzone:capped-memcry> end

D.
From the global zone, issue these commands:
zonecfg -z   dbzone
     zonecfg:dbzone> set physical=500m
     zonecfg:dbzone> end

E.
From the global zone, issue these commands:
prctl -n  zone.max -memory  -v 500M -r -i  dbzone
```

A. Option A
B. Option B
C. Option C
D. Option D
E. Option E

Answer: C

Explanation:
Add a memory cap.
zonecfg:my-zone> add capped-memory
Set the memory cap.
zonecfg:my-zone:capped-memory> set physical=50m End the memory cap specification.
zonecfg:my-zone:capped-memory> end

NEW QUESTION 92
Which five statements describe options available for installing the Oracle Solaris 11operating 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 deskto
C. The GNOME desktop package must he 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 96
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/rdsk/c4t0d0endBoot the zone and log in the zone as roo
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 97
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 [-HIaflpr] [-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 99
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 syste
B. When the system is booted, execute the zpool clear pool1 command.
C. Shut the system down, replace disk c3t3d0, and boot the syste
D. When the system is booted execute the zpool online pool1 command.
E. Shut the system down, replace disk c3t3d0, and boot the syste
F. When the system is booted, execute the zpool replace pool1 c3t3d0 command.
G. Shut the system down, replace disk c3t3d0, and boot the syste
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/rdsk/c1t0d0s0 x86# installgrub /boot/grub/stage1 /boot/grub/stage2 /dev/rdsk/c1t9d0s0

NEW QUESTION 102
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 104
Which command would you use to determine which package group is installed on your system?

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

Answer: B

Explanation:
The pkg info command provides detailed information about a particular IPS package. Note: The pkginfo command does the same for any SVR4 packages you may have
installed on the same system.
pkg info example:
$ pkg info p7zip Name: compress/p7zip
Summary: The p7zip compression and archiving utility
Description: P7zip is a unix port of the 7-Zip utility. It has support for numerous compression algorithms, including LZMA and LZMA2, as well as for various archive and compression file formats, including 7z, xz, bzip2, gzip, tar, zip (read-write) and cab, cpio, deb, lzh, rar, and rpm (read-only).
Category: System/Core State: Installed Publisher: solaris Version: 9.20.1
Build Release: 5.11
Branch: 0.175.0.0.0.2.537
Packaging Date: Wed Oct 19 09:13:22 2011
Size: 6.73 MB
FMRI: pkg://solaris/compress/p7zip@9.20.1, 5.11-0.175.0.0.0.2.537:20111019T091322Z

NEW QUESTION 106
You are executing this command in the default shell: sleep 5000 &
The system displays a number. This value is .

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

Answer: C

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

NEW QUESTION 110
You are attempting to troubleshoot an event that should have made an entry into the messages log. This event happened about two weeks ago. Which file should you look at
first?

A. /var/adm/messages
B. /var/adm/messages.0
C. /var /adm/messagas.1
D. /var/adm/messages.2

E. /var/adm/messages.3

Answer: A

Explanation:
The /var/adm/messages is the file to which all the messages printed on the console are logged to by the Operating System. This helps to track back check the console messages to troubleshoot any issues on the system.
Syslog daemon also writes to this /var/adm/messages file.
The /var/adm/messages file monitored and managed by newsyslog and its configuration file is /usr/lib/newsyslog.
This script runs as the roots cron job everyday, checks the /var/adm/messages file and copies/moves it to /var/adm/messages.0, 1, 2, 3, 4, 5, 6, 7. In other words, it does the Log Rotation for the /var/adm/messages.
In an event the /var file system is running out of space, these files needs to checked and can be removed (not the actual /var/adm/messages itself) to free up space on the file system.
However, care has to be taken, if you decide to empty the /var/adm/messages itself for any reason. This process is called Truncation.
SOLARIS SYSTEM ADMIN TIPS, /var/adm/messages

NEW QUESTION 114
You need to know the IP address configured on interface net3, and that the interface is up. Which command confirms these?

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

Answer: C

Explanation:
Show address information, either for the given addrobj or all the address objects configured on the specified interface, including the address objects that are only in the persistent configuration.
State can be: disabled, down, duplicate, inaccessible, ok, tentative Example:
ipadm show-addr
ADDROBJ TYPE STATE ADDR
lo0/v4 static ok 127.0.0.1/8 lo0/v6 static ok ::1/128

NEW QUESTION 115
.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

1Z0-821 Practice Exam Features:

* 1Z0-821 Questions and Answers Updated Frequently

* 1Z0-821 Practice Questions Verified by Expert Senior Certified Staff

* 1Z0-821 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* 1Z0-821 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year