



Oracle

Exam 1z0-822

Oracle Solaris 11 Advanced System Administration

Version: 8.5

[Total Questions: 140]

Question No : 1

You need to make a permanent modification to a machine's naming service. You need to change naming service from files to LDAP. Which two alternative methods would accomplish this task?

A. Modify the existing `/etc/nsswitch.conf` file and execute:

```
# nscfg import -f name-service/switch
```

B. Modify the existing `/etc/nsswitch.conf` file and execute:

```
# netcfg import name-service/switch
```

```
# svcadm refresh name-service/switch
```

C. Modify the existing `/etc/nsswitch.conf` file and execute:

```
# svcadm restart milestone/self-assembly
```

D. Use the `netcfg` command to modify the location profile for the DefaultFixad network configuration profile.

E. Use the `svccfg` command to set `config/*` properties in the `name-service/switch` service:

```
# svcadm refresh name-service/switch
```

```
# svcadm restart name-service/switch
```

Answer: A,D

Explanation: A:

* `nscfg`

- import, export name service configurations

* subcommand import

import [-fvq] FMRI

If none of the SMF repository properties for the specified FMRI are currently populated, import the legacy configuration files associated with the specified FMRI into the SMF repository.

* Example: Importing DNS Client Configuration

The following command imports the DNS client configuration, stored in `resolv.conf`, into the SMF repository.

```
# nscfg import svc:/network/dns/client:default
```

* `svcadm`– manipulate service instances

D: `netcfg`

- create and modify network configuration profiles

The `netcfg` utility manipulates system network configuration profiles. `netcfg` can be invoked

interactively, with an individual subcommand, or by specifying a command file that contains a series of subcommands.

/ netcfg location properties include:

/ nameservices: enum value list: files | dns | nis | ldap

Specifies the name services that should be configured, such as DNS, NIS, and LDAP.

Incorrect:

Not B: netcfg has no subcommand import.

Note:

* Each workstation has a nsswitch.conf file in its /etc directory. Each line of that file identifies a particular type of network information, such as host, password, and group, followed by one or more sources, such as NIS+ tables, NIS maps, the DNS hosts table, or local /etc, where the client is to look for that information.

* The Solaris platform provides the following naming services.

/ DNS, the Domain Name System

/ /etc files, the original UNIX naming system

/ NIS, the Network Information Service

/ NIS+, the Network Information Service Plus

/ LDAP, the Lightweight Directory Access Protocol

Question No : 2

You added an interface to the IPMP group ipmp0 with the following commands:

```
# ipadm create-ip net4
```

```
# ipadm add-ipmp -i net4 ipmp0
```

Which statement regarding the newly added interface is correct?

A. Link state failure detection will not occur.

B. Probe-based failure detection will not occur.

- C. Without a data IP address, it cannot become active.
- D. These changes will be lost when the system is booted.

Answer: B

Question No : 3

What is the purpose of the Service Management Facility (SMF) profiles?

- A. an XML file that describes current services and the instances
- B. allows the customization of services and instances
- C. stores configuration information about each service instance
- D. used to start and stop processes or services

Answer: B

Explanation: Profiles are, in many ways, similar to manifests in that they use the same XML DTD. However, instead of providing information about a service, its dependencies, and methods, a profile is used to provide customization of a service or an instance of a service. Customizations include whether an instance of a service should be enabled or disabled and any modifications to service configuration properties.

Incorrect:

Not A: Manifests (not profiles) are used to describe services and instances of a service, including any property groups and properties they might have.

Question No : 4

Which three statements describe Solaris 11 boot environments (BEs)?

- A. A full backup of your OS image is provided.
- B. Packages can be installed and uninstalled in an inactive BE.
- C. The OS can be upgraded in an active BE while the system is live without impacting production.
- D. A new BE can be created from the snapshot of an existing BE.
- E. A BE can become active without rebooting the system.
- F. An active BE can be unmounted and upgraded without impacting production.

Answer: A,C,D

Explanation:

The beadm utility enables you to create a new boot environment based on an existing snapshot.

Note:

* You can manage the boot environments on your system either by using the beadm command or by using the Package Manager.

* 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.

Question No : 5

You are the primary administrator for a set of Oracle Solaris 11 servers. You noticed some changes to configuration files. You are concerned that someone may have unauthorized access and that an authorized user may be abusing the access privilege. You want to track users of these systems to determine what tasks each user performs. Select the best way to gather this information.

- A. Solaris auditing
- B. the system/event service
- C. the system-logging service
- D. Basic Audit Reporting Tool
- E. System Extended Accounting

Answer: A

Explanation: Solaris auditing keeps a record of how the system is being used. The audit service includes tools to assist with the analysis of the auditing data.

Incorrect:

not C: Basic Audit Reporting Tool

BART is a file tracking tool that operates entirely at the file system level. Using BART gives

you the ability to quickly, easily, and reliably gather information about the components of the software stack that is installed on deployed systems. Using BART can greatly reduce the costs of administering a network of systems by simplifying time-consuming administrative tasks.

Note:

* The audit service makes the following possible:

Monitoring security-relevant events that take place on the host

Recording the events in a network-wide audit trail

Detecting misuse or unauthorized activity

Reviewing patterns of access and the access histories of individuals and objects



Discovering attempts to bypass the protection mechanisms

Discovering extended use of privilege that occurs when a user changes identity

* Auditing is the collecting of data about the use of system resources. The audit data provides a record of security-related system events. This data can then be used to assign responsibility for actions that take place on a host. Successful auditing starts with two security features: identification and authentication. At each login, after a user supplies a user name and password, a unique audit session ID is generated and associated with the user's process. The audit session ID is inherited by every process that is started during the login session. Even if a user changes identity within a single session, all user actions are tracked with the same audit session ID.

Question No : 6

You are creating a new SMF service named newservice. You perform the following steps:

-  Create the XML manifest file to define the service.
-  Create a script to be used to start and stop the service and set the execute permissions on this script.

What is the next step that you must perform to install this service?

- A. Enable the service
- B. Export the service
- C. Import the service.
- D. Create a snapshot of the service to be stored in the repository.

Answer: C**Explanation:** Refer to step 3 below.

After a typical software installation, there can be a half dozen or more processes that need to be started and stopped during system startup and shutdown. In addition, these processes may depend on each other and may need to be monitored and restarted if they fail. For each process, these are the logical steps that need to be done to incorporate these as services in SMF:

1. Create a service manifest file.
2. Create a methods script file to define the start, stop, and restart methods for the service.
3. Validate and import the service manifest using `svccfg(1M)`.
4. Enable or start the service using `svcadm(1M)`.
5. Verify the service is running using `svcs(1)`.

Question No : 7

Given the following commands and output:

```

root@Sol11-server:~# beadm list
BE          Active Mountpoint Space   Policy Created
-----
solaris NR  /           4.91G  static 2011-11-09 12:44

root@Sol11-server:~# beadm create solaris-2

root@Sol11-server:~# zlogin ozone beadm list
BE          Active Mountpoint Space   Policy Created
-----
ozone2     -           -      18.52M  static 2011-12-18 23:08
solaris    NR  /           425.87M static 2011-11-22 09:23
solaris-1 !R  -           3.0K   static 2011-12-18 23:43

```

Which statement summarizes this sequence of commands?

- A. The BE `solaris` in the Global zone is the same BE that is listed in the third command.
- B. The BE `solaris-1` is activated to boot when the `ozone` zone is rebooted.
- C. The BE `solaris-1` was created when the BE `solaris-2` was created.
- D. The BE `ozone2` is a BE of a zone with the zone name `ozone2`.

Answer: A**Explanation:**

Incorrect:

Not B: You cannot activate an unbootable BE in a nested BE.

Not D: The zone is named ozone (not ozone2).

Note:

* beadm supports the concept of a nested BE, specifically, as it pertains to BEs for non-global zones.

* beadm list [-a | -ds] [-H] [beName]

Lists information about the existing boot environment named beName, or lists information for all boot environments if beName is not provided. The Active field indicates whether the boot environment is active now, represented by N; active on reboot, represented by R; or both, represented by NR. Unbootable BEs inside of a nested BE are represented by an exclamation point (!)

Question No : 8

Examine the following command:

```
ipadm create-addr -T static -a 192.168.1.112/26 net0/v6
```

Which two statements are true?

- A. The interface is plumbed.
- B. The interface is marked down.
- C. The netmask value is fffffc0.
- D. The link local IPv6 address fe80::112 is created.
- E. Multicast datagrams are not enabled on this interface.

Answer: A,C

Question No : 9

Which vmstat field output indicates the effort the system is making to find memory for future processes?

- A. re
- B. sr
- C. free
- D. fr

Answer: A

Explanation: re

page reclaims.

The information is given in units per second.

Note:

* vmstat reports virtual memory statistics regarding kernel thread, virtual memory, disk, trap, and CPU activity.

Incorrect:

Not B: sr, pages scanned by clock algorithm

Not C: free

size of the free list (Kbytes)

Not D: fr

kilobytes freed

Question No : 10

Which two statements describe the capabilities of the Distribution Constructor?

- A. ISO images for use with the Automated Installer (AI) can be created.
- B. Bootable USB images can be created for SPARC and x86 architectures.
- C. A single installation server can be used to create ISO images, for SPARC and x86 architectures.
- D. Checkpoints are used to pause the build, thereby allowing the running of a script to modify the resulting ISO image.
- E. A single installation server can be used to create ISO images for Solaris 10 and Solaris 11.0 operating systems.

Answer: A,D

Explanation: A: Oracle Solaris Image Types include:

x86 or SPARC ISO Image for Automated Installations – The Oracle Solaris release

includes the automated installer tool. The automated installer (AI) is used to automate the

installation of the Oracle Solaris OS on one or more SPARC and x86 systems over a network.

D:

- * You can also create custom scripts to modify your installation image. Then, you can add checkpoints to the manifest file to run these custom scripts.
- * You can use the options provided in the `distro_const` command to stop and restart the build process at various stages in the image-generation process, in order to check and debug the image that is being built. This process of stopping and restarting during the build process is called checkpointing. Checkpointing is optional. Default checkpoints are specified in each manifest file.

Incorrect:

Not B: Only for x86, not for SPARC-

Oracle Solaris x86 LiveCD – You can create an x86 ISO image that is comparable to the LiveCD image that's distributed as an Oracle Solaris release. You can also modify the content of this ISO image by adding or removing packages. You can revise the default settings for the resulting booted environment to create a custom ISO image or USB image.

Note: The distribution constructor creates images based on settings specified in XML files, called manifest files. The manifest files contain specifications for the contents and parameters for the ISO images that you create using the distribution constructor. The `distribution-creator` package provides sample manifests that can be used to create a custom x86 Live Media ISO, an x86 or SPARC Automated Install ISO image, or an x86 or SPARC text installation ISO image.

The elements in each manifest file provide preset, default values that will create the type of ISO image you need. You can manually edit these preset elements in a manifest file to customize the resulting image. In addition, you can create custom scripts to further modify your image. Then, reference the new scripts in the manifest file.

Question No : 11

Which four configuration elements are managed by the `dumpadm` utility?

A. the location of the dump device

- B. the location of the savecore directory
- C. the index number for the next core dump
- D. the reserved file system space that a core dump may not use
- E. the compression of the crash dump file
- F. the size of the dump device

Answer: A,B,D,E

Explanation: dumpadm - configure operating system crash dump

SYNOPSIS

```
/usr/sbin/dumpadm [-nuy] [-c content-type] [-d dump-device]  
[-m mink | minm | min%] [-s savecore-dir]  
[-r root-dir] [-z on | off]
```

The options include:

A: -d dump-device

Modify the dump configuration to use the specified dump device.

B: -s savecore-dir

Modify the dump configuration to use the specified directory to save files written by savecore.

D: -m mink | minm | min%

Create a minfree file in the current savecore directory indicating that savecore should maintain at least the specified amount of free space in the file system where the savecore directory is located.

E: -z on | off

Modify the dump configuration to control the operation of savecore on reboot. The options are on, to enable saving core files in a compressed format, and off, to automatically uncompress the crash dump file. The default is on, because crash dump files can be very large and require less file system space if saved in a compressed format.

Question No : 12

Which option provides limits for physical, swap, and locked memory?

- A. rcap.max-rss
- B. zone.max-locked-memory
- C. capped-memory
- D. zone.max-memory

Answer: C

Explanation: The capped-memory resource sets limits for physical, swap, and locked memory. Each limit is optional, but at least one must be set.

Question No : 13

Which utility/service must you use to set processes with FSS by default?

- A. priocntl
- B. svc:/system/scheduler:default
- C. dispadmin
- D. projmod

Answer: C

Explanation: Set the default scheduler for the system to be the FSS.

```
# dispadmin -d FSS
```

Question No : 14

Which two statements regarding the pkg command are correct?

- A. It requires HTTP to connect to a remote repository.
- B. It uses the set-publisher subcommand to remove an origin.
- C. It cannot point to both sticky and nonsticky publishers in the same repository.

D. It uses the unset-publisher subcommand to remove publishers.

E. It uses the set-publisher subcommand to remove publishers.

Answer: A,B

Explanation: A: Configure pkg.depotd to provide remote access. pkg.depotd provides an HTTP interface to a pkg repo. Here we are going to make the repo server listen on port 10000, and use the repo dir we created as its default.

```
# svcadm disable pkg/server
```

```
# svccfg -s pkg/server setprop pkg/inst_root = /data/myrepo
```

```
# svccfg -s pkg/server setprop pkg/port = 10000
```

```
# svcadm refresh pkg/server
```

```
# svcadm enable pkg/server
```

B: Set-Publisher

With -G (--remove-origin), remove the URI or path from the list of origins for the given publisher. The special value * can be used to remove all origins.

Incorrect:

Not D: unset-publisher publisher ...

Remove the configuration associated with the given publisher or publisher

Not E: set-publisher

Update an existing publisher or add a package publisher. If no options affecting search order are specified, new publishers are appended to the search order and are thus searched last.

Question No : 15

The zpool configuration on serverA is:

pool 1

c3t2d0

c3t3d0

pool 2

c3t4d0

c3t5d0

The zpool configuration on servetB is:

pool1

mirror-0

c3t2d0

c3t3d0

mirror-1

c3t4d0

c3t5d0

Which option will modify the configuration on serverA to match serverB?

- A. zpool destroy pool2zpool attach pool1 c3t4d0 c3t5d0
- B. zpool destroy pool2zpool attach pool1 c3t2d0 c3t2d0 c3t4d0 c3t5d0
- C. zpool destroy pool2zpool add pool1 c3t4d0 c3t5d0
- D. zpool destroy pool2zpool mirror pool1 pool2
- E. zpool destroy pool2zpool attach pool1 c3t2d0 attach pool1 c3t3d0zpool attach pool1 c3t4d0 attach pool1 c3t5d0
- F. zpool destroy pool1; zpool destroy pool2; zpool create pool1 mirror c3t2d0 c3t3d0 mirror c4t4d0 c3t5d0

Answer: F

Explanation: Example;root@solaris:~# zpool create pool1 mirror c8t0d0 c8t1d0 mirror c8t3d0 c8t4d0

root@solaris:~# zpool status

pool: pool1

state: ONLINE

scan: none requested

config:

NAME STATE READ WRITE CKSUM

pool1 ONLINE 0 0 0

mirror-0 ONLINE 0 0 0

c8t0d0 ONLINE 0 0 0

c8t1d0 ONLINE 0 0 0

```
mirror-1 ONLINE 0 0 0  
c8t3d0 ONLINE 0 0 0
```

Question No : 16

Consider the following commands on a newly installed system:

```
zfs set compression=on rpool
```

```
zfs get -H -o source compression rpool
```

What is the output of the second command?

- A. default
- B. -
- C. local
- D. on

Answer: C

Explanation: The `zfs get` command supports the `-H` and `-o` options, which are designed for scripting. You can use the `-H` option to omit header information and to replace white space with the Tab character. Uniform white space allows for easily parseable data. You can use the `-o` option to customize the output in the following ways:

* The literal name can be used with a comma-separated list of properties as defined in the Introducing ZFS Properties section.

* A comma-separated list of literal fields, name, value, property, and source, to be output followed by a space and an argument, which is a comma-separated list of properties.

The following example shows how to retrieve a single value by using the `-H` and `-o` options of `zfs get`:

```
# zfs get -H -o value compression tank/home  
on
```

Question No : 17

To reduce the use of storage space on your server, you want to eliminate duplicate copies of data in your server's ZFS file systems. How do you specify that pool1/data should not contain duplicate data blocks on write operations?

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

Answer: E

Explanation: To create a file system with deduplication:

```
root@solaris:~# zfs create -o dedup=on
```

Note: If you have a storage pool named 'tank' and you want to use dedup, just type this:

```
zfs set dedup=on tank
```

Question No : 18

At which two stages are backups of the service configuration repository made?

- A. during boot, after the early-manifest-import service finishes
- B. during boot, before the manifest-import service starts
- C. during manifest import, before the manifest-import service finishes
- D. during manifest import, after the manifest-import service finishes
- E. when a service is modified using the svccfg command
- F. after the last successful boot
- G. when the svcadm refresh command is executed

Answer: B,D

Explanation: <http://illumos.org/msg/SMF-8000-MY>

Question No : 19

Which two zpool subcommands will permanently remove a submirror from active storage pool?

- A. remove
- B. detach
- C. destroy
- D. offline
- E. replace
- F. split
- G. zpool does not permit this operation on an active storage pool unless the submirror faults.

Answer: A,B

Explanation:

zpool detach pool device

Detaches device from a mirror. The operation is refused if there are no other valid replicas of the data.

Question No : 20

You must configure your server to use IPMP with probe based failure detection enabled. Which statement is a valid constraint or feature that applies to this requirement?

- A. Link-based detection is supported only on Generic Lan Driver version 2 (GLDv2)-complaint NICs.
- B. GLDv2 NICs are not supported in Oracle Solaris 11.
- C. GLDv3 NICs configured for link based detection by default.
- D. You must first disable link based detection before configuring probe-based failure

detection.

Answer: C

Explanation: Network drivers that support link-based failure detection monitor the interface's link state and notify the networking subsystem when that link state changes.

Incorrect:

Not B: GLDv2 is a multi-threaded, clonable, loadable kernel module that provides support to device drivers for local area networks. Local area network (LAN) device drivers in the Solaris OS are STREAMS-based drivers that use the Data Link Provider Interface (DLPI) to communicate with network protocol stacks.

Not D: Link-based failure detection is always enabled, provided that the interface supports this type of failure detection.

You cannot disable link-based failure detection if this method is supported by the NIC driver.

Note:

* To write a network driver for the Oracle Solaris OS, use the Solaris Generic LAN Driver (GLD) framework.

/ For new Ethernet drivers, use the GLDv3 framework.

/ To maintain older Ethernet, Token Ring, or FDDI drivers, use the GLDv2 framework.

* To ensure continuous availability of the network to send or receive traffic, IPMP performs failure detection on the IPMP group's underlying IP interfaces. Failed interfaces remain unusable until they are repaired. Remaining active interfaces continue to function while any existing standby interfaces are deployed as needed.

The in.mpathd daemon handles the following types of failure detection:

/ Probe-based failure detection, of two types:

No test addresses are configured (transitive probing).

Test addresses are configured.

/ Link-based failure detection, if supported by the NIC driver

Question No : 21

You are testing the connectivity between an Oracle Solaris 11 system and a local IPS server that has the host name of mercury within the domain purple.com.

The command `ping mercury` indicates the server is alive.

The URI `http://mercury.purple.com` produces the error:

Firefox can't find the server at mercury.purple.com

You enter the command `svccprop -p config network/dns/client`

Which two can be verified?

- A. the domain name of the local system
- B. the name service switch configuration
- C. the IP address of the IPS server
- D. the IP address of the DNS server
- E. the IP address of the local system
- F. the host name of the local system

Answer: A,D

Explanation: Example. Displaying Administratively Customized Properties (here only admin layer is displayed with `-l`)

The following command uses SMF layers to display administratively customized properties.

```
example% svccprop -p config -l admin svc:/network/dns/client
```

```
(A) config/domain astring admin my.domain.com
```

```
(D) config/nameserver net_address admin 10.22.33.44 10.44.33.11
```

Note:

* `config/nameserver` refers to the Solaris DNS server.

/ The `nameserver` keyword specifies DNS servers to query using IP address.

Example:

The `/etc/resolv.conf` file contains configuration directives for the DNS resolver. The following `resolv.conf` example shows two name servers and three search suffixes:

```
domain nj.bigcorporation.com
nameserver 192.168.10.11
nameserver 192.168.20.88
```

* The svcprop utility prints values of properties in the service configuration repository. Properties are selected by -p options and the operands

Question No : 22

You are configuring a system on your network that was installed using LiveCD. You configured applied a static IP address to the system. You now need to configure a default router. Assume that name services are file based and the router's IP address is 172.31.10.1.

Which two methods should you use to configure a default route on this Oracle Solaris 11 system?

- A. svccfg -s network/physical setprop config/defroute=172.31.10.1;svcadm refresh network/physical; svcadm restart network/physical
- B. svccfg -s routing/route setprop config/defroute=172.31.10.1;svcadm refresh routing/route; svcadm restart routing/route
- C. Add the IP address to the /etc/defaultrouter file and run route add default 172.31.10.1.
- D. netadm modify -p ncp DefaultFixed set-prop -p defroute=172.31.10.1
- E. ipadm set-prop -p defaultrouter=172.31.10.1

Answer: C

Explanation: * AI Server Software Requirements include:

Default router

Ensure that your AI server has a default route set by using the netstat command to show network status. If your AI server does not have a default route set, you can set a static default route by populating the /etc/defaultrouter file with the IP address of a static default route for your server's network.

* For any currently active NCP (fixed or reactive), use the route command with the -p option to persistently add a route:

```
# route -p add default ip-address
```

Because this command applies the specified route to the currently active NCP, the default route is removed and potentially replaced, if the active NCP changes.

Incorrect:

Not A, Not B: The svccfg command manipulates data in the service configuration repository. svccfg can be invoked interactively, with an individual subcommand, or by specifying a command file that contains a series of subcommands.

Question No : 23

You are about to configure an AI server and you need to determine if NWAM is configured, if the system has a manually configured IP interface. Which command gives you this information?

- A. nscfg list
- B. netadm list
- C. netcfg list
- D. svcs network/physical

Answer: B

Question No : 24

How do you add a test address to an IPMP group?

- A. Use ipadm create-addr to add the address to a member interface.
- B. Use ipadm create-addr to add the address to the IPMP interface.
- C. Use ipadm create-ip to add an address to a member interface.
- D. Use ipadm create-ip to add an address to the IPMP interface.

Answer: A

Question No : 25

You decide to create a new rights profile to include a selection of Solaris authorizations and commands. The commands in your selection will require extra privileges.

Which two files will you modify to add these privileges and authorizations?

- A. /etc/user_attr
- B. /etc/security/auth_attr
- C. /etc/security/prof_attr
- D. /etc/security/exec_attr
- E. /etc/security/prof_attr.d/core-os
- F. /etc/security/auth_attr.d/core-os

Answer: A,C

Explanation: Rights profiles – The user_attr, prof_attr, and exec_attr databases are now read-only. These local files databases are assembled from fragments that are located in /etc/user_attr.d, /etc/security/prof_attr.d, and /etc/security/exec_attr.d.

The fragment files are not merged into a single version of the file, but left as fragments. This change enables packages to deliver complete or partial RBAC profiles. Entries that are added to the local files repository with the useradd and profiles commands are added to the local-entries file in the fragment directory.

Question No : 26

The zfs holds command displays the following information:

NAME	TAG	TIMESTAMP
pool12/data@nov keep	Wed May 30	12:15:12 2012

Which two statements are true?

- A. Use zfs destroy -d pool12/data@nov to destroy the snapshot immediately.
- B. Attempts to destroy the snapshot using zfs destroy pool12/data@nov will fail.
- C. Attempts to destroy the pool12/data@nov snapshot will not destroy the snapshot immediately.

- D. The zfs directory `-R pool12/data` command will destroy the file system immediately.
- E. The `defer_destroy` property is set to on for the `pool12/data@nov` data set.
- F. The `userrefs` property is set to 1 (or higher) for the `pool12/data@nov` data set.

Answer: B,E

Explanation: * Use the `zfs holds` command to display a list of held snapshots.

* Holding a snapshot prevents it from being destroyed (B). In addition, this feature allows a snapshot with clones to be deleted pending the removal of the last clone by using the `zfs destroy -d` command. Each snapshot has an associated user-reference count, which is initialized to zero. This count increases by one whenever a hold is put on a snapshot and decreases by one whenever a hold is released.

Note:

* Example:

```
# zfs holds tank/home@now
NAME TAG TIMESTAMP
tank/home@now keep Thu Jul 15 11:25:39 2010
```

* You can use the `zfs release` command to release a hold on a snapshot or set of snapshots.

If the snapshot is released, the snapshot can be destroyed by using the `zfs destroy` command.

* Two new properties identify snapshot hold information:

The `defer_destroy` property is on if the snapshot has been marked for deferred destruction by using the `zfs destroy -d` command. Otherwise, the property is off.

The `userrefs` property is set to the number of holds on this snapshot, also referred to as the user-reference count.

Question No : 27

Your organization uses a fixed base configuration for all Oracle Solaris native brand zones that are created. You want to configure your server so that it will use your company template when the `create` command is issued. Identify the preferred way to accomplish this.

- A. Change the /etc/zones/SYSblank.xml link to link to your company template and use create -b.
- B. Set the default_template property in the system/zones service to your company template.
- C. Change the /etc/zones/SYSsolaris.xml link to link to your company template.
- D. Set the zone_default_template parameter in the /etc/default/zones file.

Answer: B

Explanation: create uses a default template of SYSdefault. The default template can be changed on a system-wide basis using the default_template SMF property of the svc:/system/zones:default service.

Note:

```
create [-F] [ -a path |-b | -t template]
```

Create an in-memory configuration for the specified zone. Use create to begin to configure a new zone

Question No : 28

By default, which directory does the audit_binfile plug-in write log files in, and what is the maximum size of each log file?

- A. /var/adm/audit, 100 MB
- B. /var/audit, 100 MB
- C. /var/adm, no limit
- D. /var/audit, no limit
- E. /var/adm/audit, 16 EB

Answer: D

Explanation: * The following directives cause audit_binfile.so to be loaded, specify the directories for writing audit logs, and specify the percentage of required free space per directory.

```
auditconfig -setplugin audit_binfile active \  
"p_dir=/var/audit/jedgar/eggplant,/var/audit/jedgar.aux/eggplant,  
/var/audit/global/eggplant;p_minfree=20;p_fsize=4.5GB"
```

* The attributes specifying the configuration of audit_binfile plugin include:

p_dir

dir1[,dir2],.. [,dirn]

A list of directories, where the audit files will be created. Any valid writable directory can be specified.

p_fsize

The p_fsize attribute defines the maximum size that an audit file can become before it is automatically closed and a new audit file is opened. This is equivalent to an administrator issuing an audit -ncommand when the audit file size equals the value specified by the administrator. The default size is zero (0), which allows the file to grow without bound.

Question No : 29

You capped the physical memory for the testzone at 50M. Which option would temporarily increase the cap to 100M?

- A. rctladm -z testzone zone.capped-memory=100M
- B. rcapadm -z testzone -m 100M
- C. rcapadm -z testzone zone.capped-memory=100M
- D. prctl testzone -m 100M

Answer: B

Explanation: How to Specify a Temporary Resource Cap for a Zone

This procedure is used to allocate the maximum amount of memory that can be consumed by a specified zone. This value lasts only until the next reboot. To set a persistent cap, use the zonecfg command.

1. Become superuser, or assume a role that includes the Process Management profile. The System Administrator role includes the Process Management profile.

2. Set a maximum memory value of 512 Mbytes for the zone my-zone.

```
# rcapadm -z testzone -m 512M
```

Question No : 30

You are using AI to install a new operating system. You add the following information to the AI manifest:

```
<configuration type="zone" name="dbzone  
source="http://sysA.example.com/zone_cfg/zone.cfg"/>
```

Which statement is true regarding the zone.cfg file?

- A. It is a text file in a zonecfg configuration format.
- B. It is an AI manifest that specifies how the zone is to be installed.
- C. It is an XML file in a form suitable for use as a command script file for the zonecfg command.
- D. It is an profile with keywords that are specific for configuring a zone as part of the installation
- E. It is an XML file that specifies the zonename, zonepath, and other zonecfg parameters.

Answer: A

Question No : 31

Your company requires all nonglobal zones to have a limit on the number of processes. The policy is designed to prevent runaway processes from impacting the global zone and other nonglobal zones. To set a limit on processes, you set the following controls:

rctl:

name: zone.max-processes

value: (priv=privileged,limit=1000,action=none)

rctl:

name: zone.max-lwps

value: (priv=privileged,limit=5000,action=deny)

You must now enable system logging for the zone.max-processes resource control to record when the limit is exceeded. Which command enables system logging for a resource

control?

- A. priocntl
- B. zonecfg
- C. rctladm
- D. prctl

Answer: C

Explanation: Global Actions on Resource Control Values

Global actions apply to resource control values for every resource control on the system. You can use the rctladm command to perform the following actions:

- * Display the global state of active system resource controls
- * Set global logging actions

Question No : 32

Your colleague is administering the company's Automated Install server and is using a custom manifest. You now plan to replace the default AI manifest with the custom manifest.

Which installadm subcommand must you use to replace the default AI manifest?

- A. update-service
- B. update-manifest
- C. set-service
- D. create-manifest
- E. create-service
- F. set-manifest

Answer: B

Explanation: Explain:

If you want to change the content of a manifest or script that has already been added to an install service, use the installadm update-manifest command. Criteria, default status, and manifest_or_script_name are not changed as a result of the update.

```
# installadm update-manifest -n s11-x86  
-f ./newregion1.xml -m region1
```

The create-manifest and update-manifest subcommands validate XML manifest files before

adding them to the install service. AI syntactically validates the AI manifests at client installation time.

Note - If an invalid manifest is provided to a client, the automated installation aborts. To investigate the cause of the validation failure, see the /system/volatile/install_log on the client.

Reference: http://docs.oracle.com/cd/E23824_01/html/E21798/changeai.html

Question No : 33

You and an associate named Frank administer a Solaris server. Frank will be responsible for monitoring system resource usage, so you set up a project entry for his account. Consider the following commands and output:

```
root@soll1-server:~# projects frank
default group.staff clockfix user.frank

root@soll1-server:~# getent user_attr frank
frank:::type=normal;project=clockfix;profiles=Audit Review,Cron Management

root@soll1-server:~# cat /etc/project
system:0:::
user.root:1:::
nopproject:2:::
default:3:::
group.staff:10:::
clockfix:510:project-clock:frank::project.cpu-caps=(priv,1.5,deny)
user.frank:501:monitor:frank::task.max-lwps=(priv,100,signal=SIGTERM)
```

Which is Frank's default project?

- A. default
- B. clockfix
- C. user.frank
- D. group.staff

Answer: A

Question No : 34

You created a virtual network of three zones.

One network hosts a web server.

Another hosts an application server used by the web server.

The third zone host a video streaming application.

You already configured a flow to prioritize the video traffic over the web server traffic. You now need to continuously monitor the flow.

Which tool must you use to gather the flow data?

- A. the system activity reporter (SAR)
- B. extended accounting
- C. the flowstat command
- D. the kstat utility

Answer: C

Explanation: Gathering Statistics About Network Traffic on Flows

Flow statistics help you evaluate packet traffic on any defined flows on the system. To obtain flow information, you use the flowstat command.

* Display statistics about incoming and outgoing packets on all flows.

flowstat

This command provides a static display of traffic information on all configured flows.

Incorrect:

Not A: In computing, sar (System Activity Report) is a Solaris-derived system monitor command used to report on various system loads, including CPU activity, memory/paging, device load, network.

Question No : 35

Which network component is the default target for IPMP probe-based Failure detection?

- A. the default router
- B. any nongateway system on the same subnet

- C. any group interface with a test address
- D. the first responder to the in.mpathd broadcast request

Answer: A

Explanation: How to Manually Specify Target Systems for Probe-Based Failure Detection

1. Add a route to a particular host to be used as a target in probe-based failure detection.

```
$ route -p add -host destination-IP gateway-IP -static
```

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

For example, you would type the following to specify the target system 192.168.10.137,

which is on the same subnet as the interfaces in IPMP group itops0:

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

This new route will be automatically configured every time the system is restarted. If you want to define only a temporary route to a target system for probe-based failure detection, then do not use the -p option.

2. Add routes to additional hosts on the network to be used as target systems.

Question No : 36

You created a role. The role should be able to change the configuration of a zone. How will you assign that privilege to the role?

- A. Modify the zone using the admin resource, set the user property to the role and the auths property to manage.
- B. Assign to the role the zone Management Rights Profile.
- C. Assign to the role the solaris.zones.* authorization.
- D. Assign to the role the zone Security Rights Profile.

Answer: B

Explanation: The Zone Management profile grants the power to manage all of the non-global zones on the system to a user.

Question No : 37

What is true about crash dump configuration?

- A. The minfree value can be to protect a percentage of available disk space.
- B. The default size of the dump device is configurable.
- C. You can use one ZFS volume for both swap and dump.
- D. You can set quota on a ZFS dump device.
- E. When set on the dump device, the minfree value sets the total size of the dump device to be a percentage of the total size of the root pool.

Answer: A

Explanation: See % below.

coredump parameter: -m mink | minm | min%

Create a minfree file in the current savecore directory indicating that savecore should maintain at least the specified amount of free space in the file system where the savecore directory is located. The min argument can be one of the following:

k

A positive integer suffixed with the unit k specifying kilobytes.

m

A positive integer suffixed with the unit m specifying megabytes.

%

A % symbol, indicating that the minfree value should be computed as the specified percentage of the total current size of the file system containing the savecore directory.

The savecore command will consult the minfree file, if present, prior to writing the dump files. If the size of these files would decrease the amount of free disk space below the minfree threshold, no dump files are written and an error message is logged. The administrator should immediately clean up the savecore directory to provide adequate free space, and re-execute the savecore command manually. The administrator can also specify an alternate directory on the savecore command-line.

Incorrect:

Not C: Separate ZFS volumes must be used for the swap area and dump devices.

Question No : 38

Which command reports the scheduling class a process falls under?

- A. ps
- B. priocnt1
- C. dispadmin
- D. rctladm
- E. prstat

Answer: A

Question No : 39

Your company decides to store its software packages in one repository. The new repository will include five publishers. How can you configure your pkg clients such that each package is updated from the publisher originally used to install it?

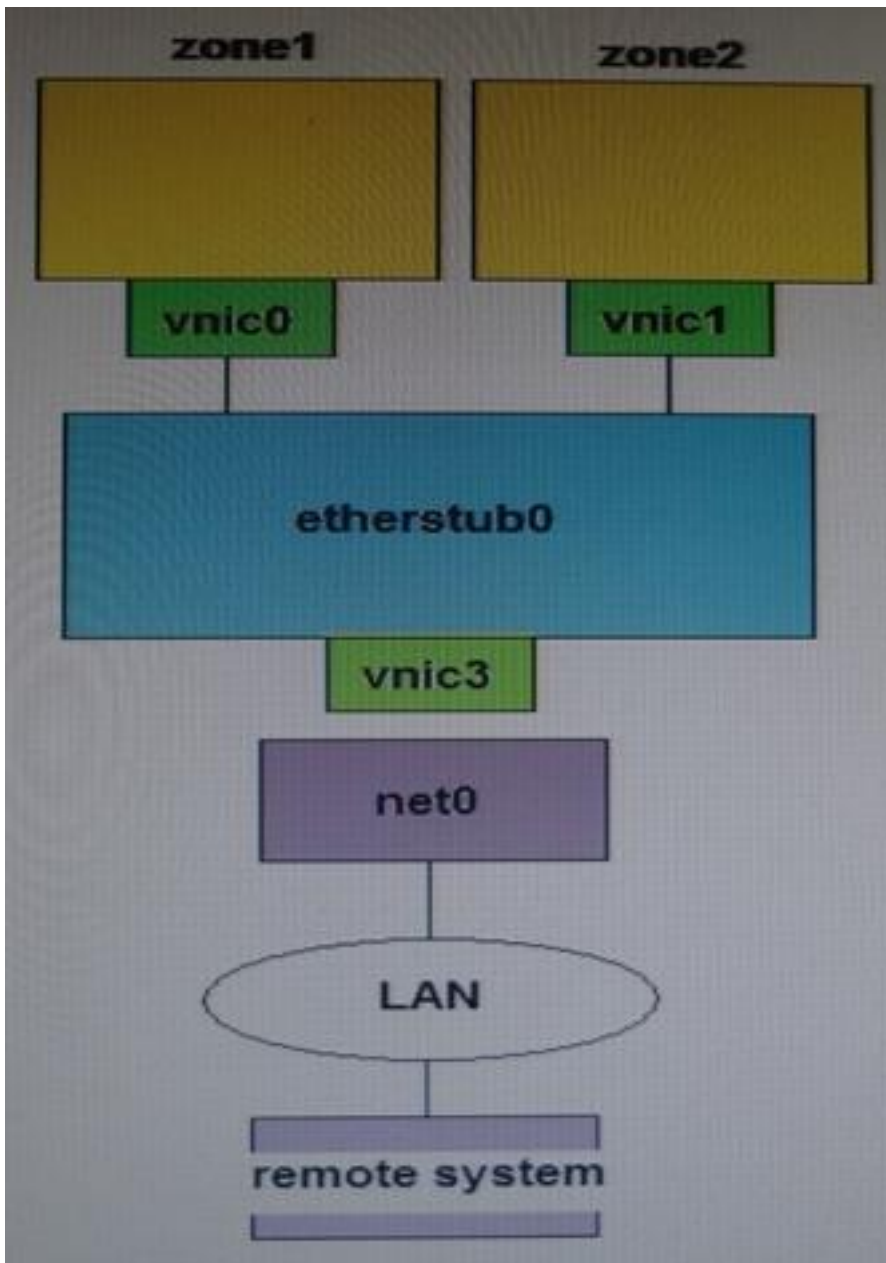
- A. Configure it such that the preferred publisher has the highest rank.
- B. Configure it such that each publisher is added with the --search-first option.
- C. Configure it such that the top ranked publisher is sticky.
- D. Configure it such that all publishers are sticky.
- E. Configure it such that the bottom-ranked publisher is sticky.

Answer: D

Explanation: set-publisher --sticky

With --sticky, specify that updates to packages that were installed from this publisher must also come from this publisher. This is the default behavior.

Question No : 40



Consider the following commands and output on the local server:

```
root@sol11-server:~# dladm show-phys
LINK          MEDIA      STATE  SPEED  DUPLEX  DEVICE
net0          Ethernet  up    1000   full    nge0
net1          Ethernet  up    100    full    nge1

zone1/etherstub0 vnic0 - 192.168.2.210
zone2/etherstub0 vnic1 - 192.168.2.220
etherstub0 vnic3 - 192.168.2.200
net0 - 192.168.1.200
remote-system 192.168.1.1
```

Also, consider the following route table entry on a remote system:

```
192.168.2.0 192.168.1.200 UG 1 1
```

You must configure a virtual switch to connect over net0 to the remote system. Select two

commands that complete the configuration.

- A. ipadm set-ifprop -p forwarding=on net0
- B. ipadm set-prop -p forwarding=on vnic2
- C. ipadm set-prop -p forwarding=on ipv4
- D. ipadm set-prop -p routing=on net0
- E. routeadm -ue ipv4-forwarding
- F. routeadm -ue ipv4-routing

Answer: A,F

Explanation:

URL: <http://www.oracle.com/technetwork/articles/servers-storage-admin/o11-118-s11-script-zones-524499.html>

Question No : 41

Your organization uses NFS to share data from Oracle Solaris servers to Oracle Solaris clients. The server currently has an NFS share configured for the rpool/export/data file system. A client is currently mounting this file system. As administrator, you add a new file system rpool/export/data/yesterdays_data, and copy data into the new file system.

Which action is required for the client currently mounting the rpool/expor/data file system to access the new data?

- A. The rpool/export/data/yesterdays_data file system must be mounted.
- B. No action is required because the data is automatically made available.
- C. The rpool/export/data file system must be remounted.
- D. The nfs/client server must be restarted.

Answer: B

Explanation:

https://docs.oracle.com/cd/E23824_01/html/821-1454/rfsadmin-56.html

Question No : 42

Which scheduling class distributes CPU resources among its processes based on assigned importance?

- A. Fair Share Scheduler (FSS)
- B. Real-Time (RT)
- C. Fixed-priority (FX)
- D. Timesharing (TS)

Answer: C

Explanation: The FX scheduler provides a scheduling policy for processes that require user or application control of scheduling priorities. The priorities of processes that run under FX are fixed.

The FX class provides a fixed-priority preemptive scheduling policy. This policy is used by processes that require user or application control of scheduling priorities but are not dynamically adjusted by the system. By default, the FX class has the same priority range as the TS, IA, and FSS classes. The FX class allows user or application control of scheduling priorities through user priority values assigned to processes within the class. These user priority values determine the scheduling priority of a fixed-priority process relative to other processes within its class.

Incorrect:

Not A: The fair share scheduling class enables you to allocate CPU time based on shares instead of the priority scheme of the timesharing (TS) scheduling class.

Not D: The goal of the time-sharing policy is to provide good response time to interactive processes and good throughput to CPU-bound processes. The scheduler switches CPU allocation often enough to provide good response time, but not so often that the system spends too much time on switching. Time slices are typically a few hundred milliseconds.

The time-sharing policy changes priorities dynamically and assigns time slices of different lengths.

Question No : 43

Your company's security policy prohibits access to the Internet. You already installed an instance of Oracle Solaris 11 on an M-series server for base testing. You used the text install media to install the system. You also installed a package repository on the same system.

There are 10 M-series servers that have just been installed on the local network.

Can you immediately install an AI server on your testing machine order to install Oracle Solaris 11 on these 10 servers?

- A. Yes, by using the existing Solaris 10 Jumpstart server.
- B. Yes, by using the text install media for the AI software.
- C. Yes, by using the Installed package repository.
- D. No, you must download the AI .iso image from Oracle first.
- E. No, the Solaris large-server group must be installed because it contains the AI setup tools.
- F. No, the Solaris 11 full n repository must be installed on the AI server.
- G. No, you must have a prebuilt image that was created by the distribution constructor.

Answer: D

Question No : 44

There is a valid SMF manifest located underneath the `/var/svc/manifest` directory.

Which four methods can be used to add it to the services repository?

- A. Reboot the system.
- B. Restart the early-manifest-import service.
- C. Use the `svccfg apply` command.
- D. Restart the manifest-import service.
- E. Use the `svccfg import` command.

Answer: A,C,D,E

Explanation: AD: Manifests from the standard directory trees `/lib/svc/manifest` and `/var/svc/manifest` are processed during system boot and anytime an administrator or program runs:

```
$ svcadm restart manifest-import
```

```
C: svccfg
```

apply subcommand

If the argument is a service profile or manifest, apply the configuration to the admin layer of the SMF repository. Services, instances, property groups, and properties will be created as necessary.

E: import [-V] [file | directory]

svccfg import on a file in a system-managed filesystem location (subdirectories of /lib/svc/manifest and /var/svc/manifest) invokes: svcadm restart manifest-import.

Placing your manifests in a system-managed location and invoking svcadm restart manifest-import to import them is the recommended practice.

svccfg import on files in other locations imports their properties as administrative customization into the admin layer. It is equivalent to:

svccfg apply [file | directory]

Incorrect:

not B: Manifests are processed in two different phases during boot.

The service svc:/system/early-manifest-import:default, a pseudo service, is responsible for the first manifest processing. This service processes only manifests from the /lib/svc/manifest directory tree before svc.startd(1M) initializes any services thus enabling services delivered in /lib/svc/manifest to always start with their most updated definition. Since this is a pseudo service, svcadm(1M) commands are ignored though svcs(1) can be used to observe status and get log file information.

The svc:/system/manifest-import:default service handles the second manifest processing and imports manifest files from both /lib/svc/manifest and /var/svc/manifest directory trees, in that respective order.

Question No : 45

Which two conditions must exist in order to add a new link to an aggregation?

- A. The new link appears in the output of dladm show-phys.
- B. The new link may already support an active interface.

- C. The LACP policy must be set to L4.
- D. The link may currently be in any state.
- E. The new link has the same MAC address as the existing links.

Answer: A,D

Explanation: A:

Note:

Use `dladm show-phys` to obtain information about the system's datalinks in relation to the physical NICs with which they are associated. Used without any options, the command displays information similar to the following:

```
# dladm show-phys
LINK MEDIA STATE SPEED DUPLEX DEVICE
net0 Ethernet up 100Mb full e1000g0
net1 Ethernet down 0Mb -- nge0
net2 Ethernet up 100Mb full bge0
net3 Infiniband -- 0Mb -- ibd0
```

* Your link aggregation configuration is bound by the following requirements:

/You must use the `dladm` command to configure aggregations.

/ (not B) An interface that has been created cannot become a member of an aggregation.

/All interfaces in the aggregation must run at the same speed and in full-duplex mode.

/ (not E) You must set the value for MAC addresses to "true" in the EEPROM parameter `local-mac-address?` For instructions, refer to [How to Ensure That the MAC Address of an Interface Is Unique](#).

Question No : 46

Which two statements describe projects and/or tasks?

- A. A task is a resource container for one process.
- B. Project resource controls are evaluated before task resource controls.
- C. Every user belongs to one or more projects.
- D. Every task associates a project with a process.
- E. A project is optional and not every user must belong to a project.