

Upgrade to Oracle Solaris 11 System Administrator

Oracle 1z0-820

Version Demo

Total Demo Questions: 23

Total Premium Questions: 235

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Managing Software Packages	50
Topic 2, Managing Boot Environments	6
Topic 3, Managing Services	14
Topic 4, Managing ZFS	50
Topic 5, Managing Network Configuration	32
Topic 6, Managing Oracle Solaris Zones	49
Topic 7, Managing User Accounts and System Security	21
Topic 8, Mix Questions	13
Total	235

QUESTION NO: 1

View the Exhibit and review the zpool and ZFS configuration information from your system.

Identify the correct procedure for breaking the /prod_data mirror, removing c4t0d0 and c4t2d0, and making the data on c4t0d0 and c4t2d0 accessible under the dev_data mount point.

- A. `zpool split pool1 pool2 c4t0d0 c4t2d0`
`zpool import pool2zfs set mountpoint=/dev_data pool2/prod_data`
- B. `zpool detach pool1 pool2`
`zpool attach pool2zfs set mountpoint=/dev_data pool2/prod_data`
- C. `zpool split pool1/prod_data -n pool2/dev_data`
`zfs set mountpoint=/dev_data pool2/prod_data`
- D. `zpool split pool1 pool2 c4t0d0 c4t2d0`
`zpool import pool2`

ANSWER: A

Explanation:

The correct procedure is `zpool split pool1 pool2 c4t0d0 c4t2d0`, followed by importing `pool2` and setting the copied dataset's mount point with `zfs set mountpoint=/dev_data pool2/prod_data`. The `zpool split` command is specifically intended to separate one side of a mirrored ZFS storage pool into a new pool. Specifying `c4t0d0` and `c4t2d0` selects the mirror devices that will form the new `pool2`, preserving the data and ZFS dataset structure from the original mirrored pool. The resulting pool must then be imported so that its datasets are available to the system. Because the copied `prod_data` dataset retains its original dataset name and mount-point property unless changed, explicitly setting `pool2/prod_data` to `/dev_data` satisfies the required new access path. The assignment syntax must be written as `mountpoint=/dev_data`, without spaces around the equals sign. Oracle documents `zpool split` as the operation for creating a new pool from a mirrored pool configuration; see the [zpool\(8\) manual page](#) and the [Oracle Solaris ZFS administration documentation](#).

QUESTION NO: 2

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

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

Your department's backup policy is to perform a full backup to a remote system disk on Saturday. On Sunday through Friday, you are to perform an incremental backup to the same remote system disk.

- Each incremental backup will contain only data that has been modified since the previous Saturday backup.
- The server file systems must remain available at all times and can never be taken offline,
- The backup must not only provide for the recovery of the most recent version of a file, but must also allow recovery of previous versions of a file created since Saturday's full backup.

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

A. On Saturday:

`zfs create snapshot pool1/data@sat.`
`zfs send pool1/data@sat > /remote/backup/full`
On each weekday:
Remove the previous daily snapshot.

```
zfs snapshot pool1/data@daily
zfs send -i pool1/data@sat pool1/data@daily > /remote/backup/full
```

B. On Saturday:

```
zfs create snapshot pool1/data@sat. zfs send pool1/data@sat | zfs recv remote/backup/'date +%m%d%Y' ' On each weekday:
```

Remove the previous daily snapshot. zfs create pool1/data@daily

```
zfs send -i pool1/data@sat pool1/data@daily | zfs recv remote/backup d%' ' ' '
```

C. On Saturday:

```
zfs snapshot pool1/data@sat
```

```
zfs send pool1/data@sat > /remote/backup/full On each weekday:
```

Remove the previous daily snapshot.

```
zfs snapshot pool1/data@daily
```

```
zfs send -i pool1/data@sat pool1/data@daily > /remote/backup/daily
```

D. On Saturday:

```
zfs create snapshot pool1/data@sat zfs send pool1/data@sat | zfs recv remote/backup On each weekday:
```

Remove the previous daily snapshot.

```
zfs create -i pool1/data@sat pool1/data@daily zfs send pool1/data@daily | zfs recv remote/backup
```

E. On Saturday:

```
zfs snapshot pool1/data@sat zfs pool1/data@sat > /remote/backup/full On each weekday:
```

Remove the previous daily snapshot.

```
zfs snapshot pool1/data@daily zfs send -i pool1/data@sat pool1/data@daily > /remote/backup/full
```

ANSWER: C

Explanation:

“On Saturday:

```
zfs snapshot pool1/data@sat
```

```
zfs send pool1/data@sat > /remote/backup/full On each weekday:
```

Remove the previous daily snapshot.

```
zfs snapshot pool1/data@daily
```

zfs send -i pool1/data@sat pool1/data@daily > /remote/backup/daily” is valid because ZFS snapshots provide a consistent point-in-time image while the pool and its mounted file systems remain online. The Saturday `zfs send` produces a complete replication stream for the `pool1/data@sat` snapshot and saves it as the full-backup stream on the remote disk. Each weekday operation creates a new snapshot and uses `zfs send -i` with the Saturday snapshot as the source and the current daily snapshot as the target. This produces an incremental stream containing the changes made since Saturday, exactly matching the stated incremental-backup definition.

Saving the full stream separately as `/remote/backup/full` and each daily incremental stream as `/remote/backup/daily` preserves the material needed to restore the Saturday baseline and subsequently apply an incremental stream. Retaining the Saturday snapshot throughout the week is essential because it is the common base for every incremental send. Removing only the prior local daily snapshot does not affect the saved backup stream and allows a fresh daily snapshot to be made. ZFS documents snapshot creation and stream generation with `zfs snapshot` and `zfs send`, including incremental replication using `-i`: [Oracle Solaris zfs\(8\) documentation](#). See also Oracle’s [ZFS snapshot and replication guidance](#).

QUESTION NO: 3

The crash dump notification on your server is:

```
Dump content: kernel and current process pages
Dump device: /dev/zvol/dsk/rpool/dump (dedicated)
Savecore directory: /var/crash
Savecore enabled: no
Save compressed: on

The files in the /var/crash directory are:
bounds      vmdump.0
```

Documentation states that there should be two core files for each crash dump in the /var/crash directory named vmdump.0

Which command should you choose to display these two files?

- A. `savecore -f vmdump.0`
- B. `dumpadm uncompressed`
- C. `gunzip vmdump.0`
- D. `dumpadm -z off`

ANSWER: A

Explanation:

`savecore -f vmdump.0` is the appropriate command because `vmdump.0` is a compressed crash-dump file created when Solaris crash-dump compression is enabled. It is not an archive intended to be viewed or extracted with a general-purpose compression utility. The Solaris `savecore` utility understands the crash-dump format and processes the dump so that the normal crash analysis files are available.


The `-f` option directs `savecore` to read crash-dump data from the named file rather than from the system's currently configured dump device. When it processes `vmdump.0`, `savecore` expands and saves the dump in the standard form used for analysis, producing the corresponding kernel image and memory core files, conventionally `unix.0` and `vmcore.0`, in the crash directory. Those files can then be used with Solaris crash-analysis tools such as `mdb`.

This behavior is specifically the purpose of the `savecore -f dumpfile` syntax: recovering a dump that has been retained in a file rather than reading directly from the configured dump device. See Oracle's [savecore\(1M\) documentation](#) and the [dumpadm\(1M\) documentation](#).

QUESTION NO: 4

On SystemA, you are troubleshooting a new installation of a virtual network. Local zones cannot connect with external systems.

View the Exhibit.



Logged in to zone1 of SystemA, you can successfully ping 172.16.0.40 and 172.16.0.42, but there is no response from any systems on the 192.168.1.x network.

Similarly, logged in to zone2 of SystemA, you can successfully ping 172.16.0.40 and 172.16.0.41, but there is no response from any systems on the 192.168.1.x network.

However, logged in to the global zone of SystemA, you can successfully ping 172.16.0.41 and 172.16.0.42, as well as 192.168.1.21 and 192.168.1.22.

Identify two possible explanations. (Choose two.)

- A. forwarding=on has not been configured for SystemA and SystemB.
- B. forwarding=on has not been configured for the global zone of SystemA.
- C. There are missing entries for SystemA and SystemB in /etc/inet/hosts.
- D. forwarding=on has not been configured for zone1 and zone2 of SystemA.
- E. ARP table entries on SystemA and SystemB are incomplete for vnic1 and vnic2.
- F. SystemB and SystemC require a static route to the 172.16.0.x network via 192.168.1.20.

ANSWER: B F

Explanation:

The global zone of SystemA is the system that connects the zones' 172.16.0.x virtual network to the external 192.168.1.x network. Therefore, it must operate as an IP router. Setting `forwarding=on` in the global zone enables IP forwarding, allowing packets received from zone1 and zone2 on the 172.16.0.x side to be routed through SystemA's external interface

toward the 192.168.1.x network. The successful pings from the global zone to both networks establish that its own interfaces and direct connectivity are functioning, while forwarding is needed for traffic transiting between those interfaces.

SystemB and SystemC also need a return path for packets addressed to the 172.16.0.x subnet. A static route for 172.16.0.0/24 with next hop 192.168.1.20 directs return traffic to SystemA, which can then deliver it to the appropriate local zone. Without this route, replies from those external systems cannot reach zone1 or zone2 even when SystemA itself can reach both external hosts. Oracle's networking documentation describes configuring IP forwarding and routing in the [Oracle Solaris 11.4 Networking Guide](#); route configuration is also covered by the [route\(8\) manual page](#).

QUESTION NO: 5

Which three statements accurately describe the Automated Installation (AI) client? (Choose three.)

- A. If the AI client does not match any criteria to use a custom manifest or script, the default manifest is used.
- B. If the AI client does not match any criteria to use a custom manifest or script, the automated installation aborts.
- C. Any manifest or script in a service can be designated to be the default for that service.
- D. Only the default.xml file is used as the default AI client manifest.
- E. If a client system does not use any SC profile, then an interactive tool opens on that client at first boot after that client installation to complete the configuration of that client.
- F. If a client system does not use any SC profile, then the install server will use the default SC profile.

ANSWER: A C E

Explanation:

AI selects installation content for a client by evaluating the client's discovered system configuration against the criteria associated with manifests and scripts in the AI service. If no custom manifest or script criteria match the client, AI uses the service's default manifest. This provides a reliable fallback installation definition, so an otherwise eligible client can still proceed through automated installation without matching a more specific client, network, memory, architecture, or disk criterion.

A service can use a manifest or script as its default by associating it without selection criteria. This enables an administrator to establish general installation behavior while retaining the ability to add more narrowly targeted manifests or scripts for clients that meet particular criteria. Oracle documents the management of AI manifests, scripts, and selection criteria in the [Oracle Solaris Automated Installer administration guide](#).

System configuration is handled after software installation. When no System Configuration profile is supplied for a client, the system starts the interactive System Configuration tool on its first boot so required settings can be completed locally. This supports unattended OS deployment while allowing initial identity, network, naming-service, and related configuration to be supplied interactively when no profile was assigned. See Oracle's [Automated Installer client installation documentation](#) for the client boot and post-installation workflow.

QUESTION NO: 6

Review the zonestat command

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

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

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

C. It is a sample of dbzone 's CPU, virtual memory, and networking utilization Physical memory is excluded from the report. The sampling is taken every 10 seconds over a 24-hour period and displayed each hour.

D. It is a sample of dbzone 's memory and CPU utilization every 10 seconds over a 24-hour period Only peak virtual memory usage and CPU utilization are displayed each hour All other utilization data is eliminated.

E. It is a sample of dbzone 's physical memory usage taken every 10 seconds over a 24-hour period. Only peak memory usage is displayed each hour All other utilization data is eliminated.

ANSWER: E

Explanation:

It is a sample of dbzone 's physical memory usage taken every 10 seconds over a 24-hour period. Only peak memory usage is displayed each hour All other utilization data is eliminated. The `-z dbzone` operand limits monitoring to the non-global zone named `dbzone`. The `-r physical-memory` resource selection restricts the utilization report to physical-memory data rather than producing the standard set of resource statistics. The final positional operands specify a 10-second sampling interval, a total monitoring duration of 24 hours, and a 60-minute report interval. Thus, `zonestat` collects samples every 10 seconds and produces a report for each one-hour reporting period. The `-R high` setting requests the high, or peak, value observed during each report interval, so each hourly result represents the highest physical-memory utilization measured from that hour's samples. The `-q` flag suppresses nonessential messages, while `-p` requests parseable output formatting; neither changes the selected resource, sampling schedule, or high-value calculation. Oracle documents the `zonestat` command options, resource filters, report types, and interval/duration/report-interval operands in the [zonestat\(1\) manual page](#).

QUESTION NO: 7

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

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

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

ANSWER: E

Explanation:

`zfs create -o dedup=on pool1/data` is correct because `dedup` is a ZFS file-system property that controls block-level deduplication. Setting it to `on` while creating `pool1/data` enables synchronous deduplication for data subsequently written to that dataset. ZFS calculates a checksum for each block and uses its deduplication table to identify blocks with identical content; instead of storing another physical copy, the new reference points to the already stored block. This can reduce space consumption where data contains significant repeated blocks, such as replicated virtual-machine images or repeated backup content.

The property can be supplied through the `-o property=value` form of `zfs create`, so it is applied immediately when the file system is created. It may also be enabled later with `zfs set dedup=on pool1/data`, but enabling it does not retroactively deduplicate blocks that were already written. In production, deduplication should be enabled only after assessing expected savings and available memory, because the deduplication table requires substantial RAM and can affect write performance. Oracle documents `dedup` as a native ZFS property and describes its on-write behavior in the [zfs\(8\) manual page](#) and the [Oracle Solaris ZFS administration documentation](#).

QUESTION NO: 8

You are setting up a local IPS package repository on your Oracle Solaris11 server: solaris.example.com.

You want to point the existing local IPS publisher to the new local IPS repository located in /repo.

These are the steps that you have followed:

1. Download and rsync the contents of the Oracle Solaris11 repository ISO image to the /repo directory.
2. Configure the repository server service properties. The svcprop command display, the IPS related properties:

```
pkg/inst_root astring/repo
```

```
pkg/readonly Boolean true
```

The ls command displays the contents of the /repo directory:

```
#ls/repo
```

```
Pkg5.repository publisher
```

The svcs publisher command shows the svc: /application/pkg/server: default service is online.

The pkg publisher command shows the svc: /application/pkg/server: default service is online.

The pkg publisher command still displays:

```
PUBLISHER TYPE STATUS URI
```

```
solaris origin online http://pkg.oracle.com/solaris/release/
```

Which steps needs to be performed to set the local IPS publisher to the local IPS repository/repo?

- A. Issue the pkgrepo refresh -s command to refresh the repository.
- B. Restart the svc:/application/pkg/server:default service.
- C. pkg set-publisher command to set the new repository location.
- D. Issue the pkgrepo rebuild command to rebuild the repository.
- E. Issue the pkgrepo set command to set the new repository location.

ANSWER: C

Explanation:

The `pkg set-publisher` command is the correct mechanism because publisher configuration is maintained on the IPS client, independently of the repository-server SMF service. Although the repository content has been placed in /repo and the package server is online, the client still has the Oracle public release repository configured as the origin for the solaris publisher. The origin must therefore be changed to the URI through which the local repository is served, such as `http://solaris.example.com/` when the package server publishes /repo at its document root. A typical command is `pkg set-publisher -G '*' -M '*' -g http://solaris.example.com/ solaris`. The `-g` option adds the specified origin, while `-G '*'` and `-M '*'` remove existing origins and mirrors so that package operations use the local source rather than the public Oracle URI. This client-side publisher update is what changes the URI displayed by `pkg publisher`. Oracle documents publisher origin management with `pkg set-publisher` in the [pkg set-publisher\(8\) manual page](#) and describes configuring local IPS package repositories in the [Oracle Solaris package repository documentation](#).

QUESTION NO: 9

You need to remove the gzip software package.

Before you remove the package, you want to create a backup of the operating environment.

Which IPS command should you use?

- A. Use the `pkg remove` command with the `--require-backup-be` option
- B. Use the `beadm create` command to create a new boot environment. Use the `pkg remove` command to remove the `gzip` package from the current boot environment.
- C. Use the `luucreate` command to create a new boot environment. Use the `pkgrm` command to remove the `gzip` package from the current boot environment
- D. Use the `pkg uninstall` command with the `-backup-be-name` option

ANSWER: A

Explanation:

Use `pkg remove --require-backup-be gzip`. The IPS `pkg remove` subcommand supports boot-environment options when a package operation changes the active boot environment. The `--require-backup-be` option requires IPS to create a backup boot environment before applying the package removal. If a backup boot environment cannot be created, the package operation is not performed, which provides the requested protection before removing `gzip`. The backup boot environment preserves the pre-change state and can be selected at boot if it is necessary to return to the system state that existed before the package removal. This is preferable to relying on a manually created copy when the requirement is specifically to have the IPS package operation enforce creation of the backup. The package name can be supplied as `gzip` (or as the package FMRI if a more specific version or publisher is required). Oracle documents `--require-backup-be` as a boot-environment option for package image-modifying operations in the [pkg\(8\) manual page](#). Additional background on managing boot environments with IPS is available in the [Oracle Solaris boot environment documentation](#).

QUESTION NO: 10

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

Choose the command used to connect to testzone's console.

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

ANSWER: E

Explanation:

The command `zlogin -C testzone` connects to the virtual console of the non-global zone named `testzone`. The uppercase `-C` option instructs `zlogin` to attach to the zone console rather than start a new login session and task within the zone. This is especially useful before or during the zone boot sequence because it provides console access to observe boot messages and interact with the zone as its console becomes available. After connecting, console behavior is similar to a system console session; to disconnect without terminating the zone, use the console escape sequence, typically `~.` entered at the start of a new line. The zone must be in an appropriate state for console access, and booting can be initiated separately with `zoneadm -z testzone boot` if necessary. Oracle documents `zlogin -C zonename` specifically as the syntax for connecting to a zone console. See the [zlogin\(1\) manual page](#) and Oracle's [Solaris Zones documentation](#) for console-login and zone administration details.

QUESTION NO: 11

You upgraded your server to Oracle Solaris 11 and you imported a zpool(pool) that was created in Solaris 10. You need to create an encrypted ZFS file system in pool, but first you need to make sure that your server supports ZFS encryption.

Which four statements are true for support of ZFS encryption?

- A. The encrypted file system must have been created in Oracle Solaris 11. To encrypt a ZFS file system from a previous version of Solaris, upgrade the zpool and create a new encrypted ZFS file system in that zpool. Finally, move or copy the data from the existing ZFS file system into the new encrypted ZFS file system.
- B. If you plan to create an encrypted file system in an existing zpool, the zpool must be upgraded to ZFS version 30, minimum.
- C. ZFS encryption is integrated with the ZFS command set and no additional packages need to be installed.
- D. Encryption is supported at the dataset (file system) level.
- E. Encryption is supported at the pool level only and every file system in the pool will be encrypted.
- F. Encryption is supported at the pool level only and every file system in the pool will be encrypted.
- G. You cannot create an encrypted file system in a zpool that was created prior to Oracle Solaris 11. Create a new zpool in Solaris 11, create an encrypted ZFS file system in the new zpool, and move or copy the data from the existing file system into the new encrypted file system.

ANSWER: A B C D

Explanation:

Native ZFS encryption in Oracle Solaris 11 is implemented through the standard ZFS administrative interface, so encryption properties and key management are configured with ZFS commands rather than through a separate encryption product. An existing pool that originated on an earlier Solaris release can be used, provided that it is upgraded to ZFS pool version 30 or later, which supplies the required encryption capability. The existing unencrypted dataset itself is not converted in place: an encrypted filesystem must be newly created under Oracle Solaris 11. Data can then be migrated into that encrypted filesystem by copying or moving it from the prior dataset. Encryption is a dataset-level feature, allowing an administrator to select the filesystem or volume that will be an encryption root instead of forcing encryption across an entire pool. Descendant datasets can inherit the encryption behavior and key context of their encryption root. These capabilities make it possible to introduce encryption gradually into an imported Solaris 10 pool after the pool has been upgraded. Oracle documents the pool-version prerequisite and creation workflow in its [Managing ZFS Encryption](#) guide, and describes ZFS dataset administration in the [zfs\(1M\) manual page](#).

QUESTION NO: 12

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

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

ANSWER: B

Explanation:

The `pkg/readonly` property for the `application/pkg/server` service must be set to `false` before publishing because publishing with `pkgsend` is a modifying operation against the IPS repository. The depot server's `pkg/readonly` setting controls whether it accepts repository changes: when it is enabled, clients can retrieve package data but cannot add or

modify package content. Its default value is `true`, so a newly configured depot commonly permits read-only access until this property is explicitly changed. Set the property through SMF, for example with `svccfg -s application/pkg/server setprop pkg/readonly=false`, then refresh or restart the service so that `pkg.depotd` uses the updated configuration. Once the depot is no longer read-only, `pkgsend` can publish package manifests and payloads to the repository. Oracle documents `pkg/readonly` as the property that disables modifying operations initiated by tools such as `pkgsend`; retrieval remains available regardless. See the Oracle [pkg.depotd manual page](#) and the [pkgsend manual page](#).

QUESTION NO: 13

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

- A. You can perform a text or LiveCD installation locally or over the network.
- B. The text Installer does not install the GNOME desktop. The GNOME desktop package must be added after you have installed the operating system.
- C. The LiveCD Installation cannot be used to install multiple instances of Oracle Solaris.
- D. 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.
- E. The LiveCD Installer is for x86 platforms only.
- F. The GUI installer cannot be used to upgrade your operating system from Solaris 10.
- G. 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.
- H. The LiveCD installer can be used for SPARC or x86 platforms.

ANSWER: B C D E F


Explanation:

The text installer does not install the GNOME desktop by default, so a system installed by that method can be given the desktop environment later by installing the appropriate desktop package set. This supports installations intended for server or console-oriented use while retaining the option to add a graphical desktop afterward. The LiveCD installer is limited to x86 systems; SPARC installations use text-based installation media and methods. The LiveCD graphical installation workflow also has disk-layout limitations that make it unsuitable where an existing VTOC slice must be preserved, and it is not the installation method for deploying multiple Solaris instances. Oracle Solaris 11 installation media provides a fresh-install path rather than an in-place upgrade path from Oracle Solaris 10. Consequently, a Solaris 10 environment must be migrated or reinstalled rather than upgraded by using the GUI installer. These characteristics distinguish the supported scope of the LiveCD GUI installer from the text installer and clarify the supported transition path to Solaris 11. Oracle's installation documentation describes the media-specific installation methods and platform support in the [Oracle Solaris 11 Installation Guide](#); later documentation also retains the distinction between installation methods in [Installing Oracle Solaris 11.4 Systems](#).

QUESTION NO: 14

You have been asked to troubleshoot the initial configuration of a virtual network connecting two local zones with the outside world.

View the exhibit.



The command

`dladm create-vnic -l vswitch192.168.1 vnic1` fails with the error `dladm: invalid link name 'vswitch192.168.1'` What is the reason for this error?

- A. The name `vswitch192.168.1` is not legal.
- B. The zone must be specified with `dladm create-vnic -z zone3 vnic1`.
- C. The virtual interface must be specified with `dladm create-vnic -z zone3 vnic1`.
- D. The virtual interface must be created with `ipadm create-vnic -l switch192.168.1`.
- E. The virtual switch must be created first with `dladm create -etherstub vswitch192.168.1`.

ANSWER: E

Explanation:

The virtual switch must be created first with `dladm create -etherstub vswitch192.168.1`. The `-l` argument of `dladm create-vnic` identifies the underlying data link on which the VNIC is created. That link must already exist and can be either a physical data link or an etherstub. In this configuration, `vswitch192.168.1` is intended to be an etherstub that provides an internal Layer 2 network for the VNICs used by the zones. Because no data link with that name exists yet, `dladm` cannot use it as the VNIC's backing link and reports it as an invalid link name.

Create the etherstub first, using the correct syntax `dladm create-etherstub vswitch192.168.1`. Then create each VNIC on that etherstub, for example `dladm create-vnic -l vswitch192.168.1 vnic1`. VNICs attached to the same etherstub are connected through the implicit virtual switch associated with that etherstub. This is the standard Solaris approach when a virtual network is required without relying on a physical NIC for that internal segment. Oracle documents etherstubs and VNIC creation in the [Oracle Solaris virtual network documentation](#) and the [dladm manual page](#).

QUESTION NO: 15

You are configuring RBAC for a user account named `operator1`.

Which three statements correctly describe Oracle Solaris RBAC?

- A. Authorizations can permit security-relevant operations.
- B. Rights profiles can contain authorizations and command execution attributes.

- C. Rights profile definitions are stored in the `prof_attr` database.
- D. Roles are intended to be used for direct login at the system console.
- E. Command execution attributes are stored in the `user_attr` database.
- F. Authorizations can be assigned only to the root account.

ANSWER: A B C

Explanation:

Authorizations are rights that permit a user or role to perform security-relevant operations. They are defined in the `auth_attr` database and can be assigned to users or roles through user attributes. Rights profiles are collections of authorizations and command execution attributes, and profile definitions are maintained in the `prof_attr` database.

Roles are special accounts that are assumed by authorized users; they are not normally used for direct login. Command-specific privilege execution information belongs in the `exec_attr` database. See the [Oracle Solaris rbac\(7\) manual page](#), the [auth_attr\(5\) manual page](#), and the [prof_attr\(5\) manual page](#).

QUESTION NO: 16

New features were added to ZFS in Oracle Solaris11. Your justification to upgrade from Solaris10 to Oracle Solaris11 is that it will be possible to take advantage of the enhancements that were made to ZFS.

Identify the three ZFS functions and features that are included in Oracle Solaris 11, but not in Solaris 10.

- A. Encrypted ZFS datasets
- B. Ability for ZFS to detect and remove redundant data from the tile system
- C. Shadow Data Migration
- D. Ability to split a mirrored ZFS storage pool
- E. Ability to use ZFS on the boot drive and boot to a ZFS root file system.
- F. elimination of the swap file system when using ZFS on the root disk

ANSWER: A B E

Explanation:

Oracle Solaris 11 expands ZFS capabilities with integrated dataset encryption, deduplication, and ZFS-root boot support. **Encrypted ZFS datasets** allows encryption to be applied at the dataset level using ZFS encryption properties and keys. This provides encryption managed as part of ZFS administration, rather than relying only on an external disk-encryption layer. Oracle documents the encryption feature, including encrypted datasets and key management, in its ZFS administration guidance.

The **Ability for ZFS to detect and remove redundant data from the tile system** describes ZFS deduplication. When deduplication is enabled, ZFS identifies duplicate data blocks and retains a single stored copy while maintaining the appropriate references for files that contain the same block content. This can reduce storage consumption where duplicate blocks are common, although Oracle emphasizes that deduplication requires careful capacity planning because its deduplication tables consume memory.

The **Ability to use ZFS on the boot drive and boot to a ZFS root file system.** is also a key Solaris 11 operational capability. A ZFS root pool contains the boot environment and operating-system datasets, enabling boot environments to be managed through ZFS snapshots and clones. See Oracle's [ZFS deduplication documentation](#) and [ZFS root pool and boot documentation](#).

QUESTION NO: 17

Which three supported in previous versions of Solaris are no longer supported in Oracle Solaris 11? (Choose three.)

- A. 32-bit SPARC hardware support
- B. 32-bit x86 hardware support
- C. SPARC Enterprise M-series servers
- D. all Sun Netra series servers
- E. Oracle Solaris Containers for Linux applications

ANSWER: A B E

Explanation:

Oracle Solaris 11 ended support for 32-bit SPARC hardware support because the operating system requires a 64-bit SPARC platform. This retirement applies to older SPARC V8-class hardware; it does not prevent the use of 32-bit application binaries where their relevant compatibility support remains available on a supported 64-bit system.

Oracle Solaris 11 also ended 32-bit x86 hardware support. Solaris 11 requires a 64-bit x86 processor capable of running the 64-bit Solaris kernel. This was a platform-support change, rather than a general removal of every possible 32-bit user application environment.

Oracle Solaris Containers for Linux applications was also discontinued. This feature, commonly associated with the Linux branded-zone capability in earlier Solaris releases, is listed in Oracle's end-of-feature notices for the Solaris 11 transition. Administrators planning an upgrade therefore need to account for Linux-branded workloads by moving them to a supported virtualization, container, or native operating-system solution.

Oracle documents these removals in its [Oracle Solaris 11 End of Feature Notices](#). The Solaris 11 documentation library also provides release and administration material for supported platform and zone configurations: [Oracle Solaris 11 Documentation](#).

QUESTION NO: 18

You want to reduce the installed content in an Oracle Solaris image by disabling documentation facets.

Which three statements correctly describe IPS facets?

- A. The `pkg facet` command displays facet values configured in an image.
- B. `pkg change-facet facet.doc.man=false` can disable installation of manual-page content marked with that facet.
- C. Changing a facet can cause IPS to add or remove selected package actions without uninstalling the package itself.
- D. A facet is used to select the processor architecture of an IPS image.
- E. Facet values can be changed only when creating a new boot environment.
- F. Disabling a facet prevents all future IPS package updates.

ANSWER: A B C

Explanation:

Facets are image properties that control the installation of optional package actions, such as documentation, locale-specific files, and other selectively installed content. The `pkg facet` command displays the facets configured in the image. A facet can be changed with `pkg change-facet`; for example, `pkg change-facet facet.doc.man=false` disables installation of manual-page actions marked with that facet.

When a facet is changed, IPS updates the image as needed to add or remove package actions controlled by that facet. The underlying package does not necessarily have to be uninstalled. Facets are distinct from variants, which represent mutually

exclusive package choices such as architecture or debug versus non-debug content. See the [Oracle Solaris pkg\(8\) manual page](#).

QUESTION NO: 19

You need to remove the gzip software package.

Before you remove the package, you want to create a backup of the operating environment.

Which IPS command should you use?

- A. Use the `pkg remove` command with the `--require-backup-be` option
- B. Use the `pkg remove` command to remove the gzip package from the current boot environment.
- C. Use the `pkgrm` command to remove the gzip package from the current boot environment
- D. Use the `pkg unmsall` command with the `-backup-be-name` option

ANSWER: A

Explanation:

Use the `pkg remove` command with the `--require-backup-be` option. In Oracle Solaris IPS, package operations that alter the installed system can be performed in a new boot environment (BE), preserving the current BE as a bootable backup. The `--require-backup-be` option explicitly requires IPS to create that backup BE when needed before carrying out the package removal. If the backup BE cannot be created, the operation does not proceed, which protects the currently running operating environment from being changed without the requested recovery point.

For example, the command would be `pkg remove --require-backup-be gzip`. After the command completes, the package change is applied to the new BE, while the original environment remains available in the boot menu. This is especially useful for administrative changes because it provides a straightforward rollback path: boot the previous BE if the removal has an unexpected effect. Oracle documents boot-environment options for IPS package operations, including `--require-backup-be`, in the [pkg\(1\) manual page](#). See also Oracle's [Boot Environments and Package Updates documentation](#) for the relationship between IPS package changes and boot environments.

QUESTION NO: 20

Review the `zonestat` command:

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

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

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

ANSWER: E

Explanation:

The statement "It is a sample of dbzone's physical memory usage taken every 10 seconds over a 24-hour period. Only peak memory usage is displayed each hour. All other utilization data is eliminated." correctly describes the command. The positional operands specify a sampling interval of 10 seconds, a total duration of 24 hours, and a reporting interval of 60 minutes. Thus, zonestat collects zone statistics every 10 seconds and produces a summarized report for each one-hour interval throughout the day.

The `-z dbzone` selection limits the report to the non-global zone named dbzone. The `-r physical-memory` selection restricts output to the physical-memory resource, rather than presenting the normal collection of resource categories. The `-R high` setting requests high-water, or peak, values for the reporting interval. Consequently, each hourly result represents the highest observed physical-memory usage among the samples gathered during that hour. The repeated parsable-output setting, together with the `zones` output selection, produces machine-readable zone-oriented output; it does not alter the collection interval, duration, or peak-value semantics. Oracle documents these operands and resource/report-selection controls in the [zonestat\(8\) manual page](#) and the [Solaris Zones documentation](#).

QUESTION NO: 21

Identify two correct statements about the Distribution Constructor. (Choose two.)

- A. Custom scripts can be used to modify the resulting ISO image created.
- B. An Image Packaging System (IPS) software repository server is required during installation from an ISO image.
- C. Manifest files contain the specifications for the contents and parameters of the ISO images to be created.
- D. ISO images can be created that dual boot, supporting both x86 and SPARC architectures with a single image.
- E. ISO images for installation of a variety of Oracle Solaris versions and hardware architectures can be created by a single installation server.

ANSWER: A C

Explanation:

Custom scripts can be used to modify the resulting ISO image created because Distribution Constructor builds images through a sequence of configurable checkpoints. A manifest can include script-based customization steps, allowing an administrator to tailor the image build process and its resulting content. This supports creating installation media that incorporates site-specific configuration, packages, files, or other required image customizations as part of the automated build workflow.

Manifest files contain the specifications for the contents and parameters of the ISO images to be created because the Distribution Constructor manifest is the central definition of an image build. It identifies the image type and output characteristics and provides the instructions that determine the software content and construction process. Administrators can begin with supplied manifests and adapt them to define a custom distribution image consistently and repeatably. Oracle's Distribution Constructor documentation describes manifests as the mechanism used to specify image construction and customization, including the checkpoint workflow. See the [Oracle Solaris 11.3 Creating Custom System Images documentation](#) and the [Distribution Constructor manual page](#).

QUESTION NO: 22

You are troubleshooting the Oracle Solaris11 Automated Installer (AI), which is not connecting with the IPS software repository.

Which three steps will help determine the cause of DNS name resolution failure?

- A. Verify the contents of `/etc/resolv.conf`.
- B. Run `netstat -nr` to verify the routing to the DNS server.

- C. Ping the IP address of the IPS server to verify connectivity.
- D. On the installation server, verify that the menu.1st file for the client points to a valid boot arc hive.
- E. Run `df -k` to verify that the boot directory containing the boot archive is loopback mounted under `/etc/netboot`.
- F. Run the command `/sbin/dhccpinfo DNSserv` to ensure that the DHCP server is providing DNS server information.

ANSWER: A B F

Explanation:

DNS resolution is a prerequisite when an Automated Installer client must contact an IPS repository by hostname. Verifying the contents of `/etc/resolv.conf` establishes whether the client has resolver configuration, including usable `nameserver` entries and, where applicable, search-domain settings. Solaris uses this file to identify the DNS servers that should receive hostname lookup requests. If the file is missing or empty on a DHCP-configured client, running `/sbin/dhccpinfo DNSserv` checks the DNS-server values supplied through DHCP. This directly identifies whether the client is receiving the network configuration it needs to resolve the repository host name.

After resolver settings are present, the client must also be able to reach the configured DNS server. Running `netstat -nr` displays the routing table and allows the administrator to confirm that a default route, or a more-specific route, exists for traffic to the DNS server. Together, these checks isolate the essential DNS path: resolver configuration, DHCP delivery of DNS information when used, and network routing to the resolver. Oracle Solaris documents the resolver-file format in [resolv.conf\(5\)](#) and DHCP client information retrieval in [dhccpinfo\(8\)](#).

QUESTION NO: 23

Select the five tasks that need to be performed on the Automated Installer (AI) install server before setting up the client.

- A. Configure an accessible IPS package repository and set the publisher origin to that repository.
- B. Set up a IP address on the AI install server.
- C. The DHCP server must be enabled on the install server and must provide the DHCP service for the clients.
- D. DHCP must be available on the network for the Install server and the clients, but the install server does not need to be the DHCP server.
- E. Download the AI boot image. The image must be the same version as the Oracle Solaris OS that you plan to install on the client.
- F. Download the text install image into the IPS repository.
- G. Install the AI installation tools.

ANSWER: A B D E G

Explanation:

An AI install server must have stable network connectivity, so configuring a static IP address and a default route ensures that clients can consistently contact the service. DHCP must be available to network-booted clients because it supplies the initial network configuration and information needed to locate the installation service; it can be provided by a host other than the AI server. The server also needs the Automated Installer administration software, including `installadm`, to create and manage AI install services. An AI boot image matching the target Solaris release is required because it supplies the network boot environment used by the client. Finally, the installation workflow requires an accessible IPS package repository and a configured publisher origin, since the client obtains Solaris packages from the repository specified by its AI installation instructions. A local repository is one supported approach, although the key requirement is that the repository be reachable by the systems involved. These preparations establish the server-side network boot, service-management, boot-image, and package-delivery components before client configuration begins. Oracle's AI documentation describes the install-service and network-boot requirements in the [Oracle Solaris Automated Installer Guide](#) and the associated installation administration procedures in the [Oracle Solaris Installation Guide](#).