PassQueen
a.k.a dumpsqueen.com

Oracle Solaris 11 Advanced System Administration

Oracle 1z0-822

Version Demo

Total Demo Questions: 18

Total Premium Questions: 186

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Managing Boot and Services	31
Topic 2, Managing Software Packages and Boot Environments	33
Topic 3, Administering ZFS	22
Topic 4, Configuring and Administering Oracle Solaris 11 Zones	12
Topic 5, Administering Oracle Solaris 11 Network	45
Topic 6, Configuring Resource Management	34
Topic 7, Mix Questions	9
Total	186

QUESTION NO: 1

Which two statements are correct regarding an SMF service instance in the maintenance state?

A. svcs -xv displays diagnostic information for the failed instance.

B. svcadm clear can clear the maintenance state.

C. svcadm refresh always clears the maintenance state.

D. The service is permanently disabled until the system is rebooted.

E. SMF automatically deletes the service manifest when the instance enters maintenance.

ANSWER: A B

Explanation:

The command `svcs -xv` provides detailed diagnostic information for services in the maintenance state. Its output identifies the affected service instance, gives a reason for the failure where available, and commonly identifies the relevant service log file.

The command `svcadm clear` clears the maintenance state of an instance. If the service is enabled, SMF attempts to start it again after the maintenance state is cleared. Administrators should correct the underlying configuration or method failure before clearing the state; otherwise, the service can return to maintenance.

See the Oracle Solaris svcs(8) and svcadm(8) manual pages.

QUESTION NO: 2

You have a ZFS pool that contains a hierarchy of data file systems. You create snapshots of the file systems and you created a clone (dpool/export/CID) of the dpool/export/home/CID file systems. The file systems are as follows:

Now you remove a file from the cloned file system:

root@sll-server1:~# rm /export/CID/core.bash.8070

How will space usage be changed for dpool/export/CID?

A. The USED value will increase and the REFER value will decrease; the AVAIL value will be unchanged.

B. The USED value will decrease and the REFER value will increase; the AVAIL value will increase.

C. The USED value will decrease, the REFER value will decrease; the AVAIL value will increase.

D. USED, REFER and the AVAIL value will be unchanged.

ANSWER: A

Explanation:

The USED value will increase and the REFER value will decrease; the AVAIL value will be unchanged. A ZFS clone initially shares the blocks present in its origin snapshot, so its REFER property includes the space referenced by those shared blocks. Removing `core.bash.8070` from the clone removes the clone's reference to that file, causing the clone's REFER value to decrease.

The file's blocks cannot be released to the pool merely because they were removed from the clone: they are still retained by the origin snapshot on which the clone depends. ZFS space accounting charges this retained, no-longer-referenced-by-the-clone data to the clone's USED space. Consequently, USED increases even though the operation is a deletion rather than a write. No pool blocks are newly allocated or returned to free space by this operation, so the amount available to the clone, shown by AVAIL, does not change.

This behavior follows ZFS copy-on-write snapshot and clone semantics: snapshots preserve referenced blocks, while clones share their origin's data until changes, including deletions, alter the clone's block references. Oracle documents ZFS snapshot and clone behavior in the Oracle Solaris ZFS Administration Guide and ZFS property reporting in the ZFS space accounting documentation.

QUESTION NO: 3

A recursive snapshot was taken of the root pool and the snapshot streams are stored on a remote system. The boot disk has failed, has been replaced, and the root pool snapshots have been restored. Which two steps are still required to make the system bootable?

A. Re-create the swap and dump devices.

B. Install the boot blocks on the new disk.

C. Restore the snapshot stream.

D. Set the bootfs property on the root pool.

E. Perform a ZFS rollback to restore the file systems in the root pool.

ANSWER: B D

Explanation:

After the restored snapshot streams have recreated the root-pool datasets on the replacement disk, the pool must be configured to identify the boot environment that firmware and the boot loader should use. Setting the `bootfs` property on the root pool designates the specific ZFS dataset containing the bootable root file system. For example, the property is set with a command such as `zpool set bootfs=rpool/ROOT/solaris rpool`, using the restored boot-environment dataset name. This is necessary because the pool-level boot target must be explicitly available after rebuilding the pool from received streams.

The replacement disk also requires boot blocks to be installed. Boot blocks provide the platform boot code needed to begin loading Solaris from the ZFS root pool. On SPARC systems this is typically performed with `installboot`; on x86 systems, the relevant boot code is installed for the disk's partitioning scheme. With both the bootfs target configured and the boot blocks installed, the restored root pool contains a bootable environment and the new disk has the required code to start it. Oracle's documented root-pool recovery procedure explicitly lists setting the bootfs property and installing boot blocks as the final actions before rebooting. See Oracle Solaris ZFS Administration Guide: Re-creating a ZFS Root Pool and Oracle Solaris bootfs property documentation.

QUESTION NO: 4

The following information describes the current dump configuration of your server:

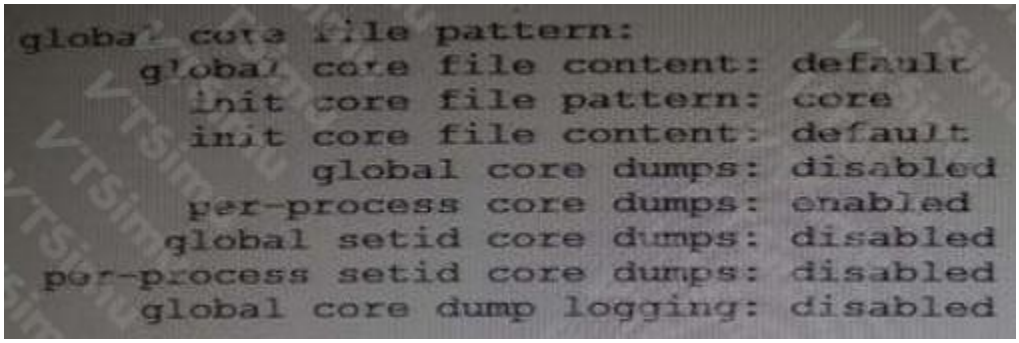

Assume that the following command lines have been run on a system:

(root) # coreadm –g $HOME/var/core/%m.core.%f.%t (non-root) $ coreadm –p core.%f.%p

Identify the result of a subsequent process crash.

A. A copy of a core file will appear in /var/core.

B. A copy of a core file will appear in the process' current working directory.

C. The root copy will include the taskid of the process.

D. The nonroot copy will include the day of the process crash.

E. The global format will override the per-process format.

F. The global core dump is disabled, so a core file will not be saved.

ANSWER: B

Explanation:

A copy of a core file will appear in the process' current working directory. The non-root invocation of `coreadm -p core.%f.%p` establishes the per-process core-file naming pattern. Because that pattern is relative rather than an absolute pathname, Solaris creates the per-process core file relative to the crashing process's current working directory. The `%f` token is expanded to the executable file name and `%p` to the process ID, so the resulting file name has the form `core.executable-name.process-id`.

Per-process core patterns are inherited by child processes, which is why this setting controls a subsequent crash in the user's process tree. A global core pattern configured by root is a separate system-wide output mechanism; it does not replace the applicable per-process pattern. The configured global pathname also uses `$HOME` as expanded by the root shell and includes the machine-name, executable-name, and time substitutions. Oracle documents that relative core-file patterns are interpreted relative to the process current directory, and that `-p` configures the pattern for a process and its descendants. See the coreadm(1M) manual page and the core(5) manual page.

QUESTION NO: 5

Which two statements are correct regarding Solaris process privilege sets?

A. The effective set contains privileges currently in effect for a process.

B. The limit set defines the maximum privileges a process and its descendants can obtain.

C. All processes in the global zone automatically have every Solaris privilege in their effective set.

D. A process can acquire any privilege that is not present in its limit set.

E. The permitted set is used only for file-system access checks.

ANSWER: A B

Explanation:

The effective privilege set contains the privileges that are currently in effect for privilege checks made by the process. A privilege must be effective when the process performs an operation that requires that privilege, subject to the normal Solaris privilege model.

The limit privilege set establishes the maximum set of privileges that a process and its descendants can obtain. A process cannot acquire a privilege outside its limit set. Administrators can use limit sets, including zone privilege limits, to ensure that a workload cannot obtain excluded privileges.

See the Oracle Solaris privileges(7) manual page and the ppriv(8) manual page.

QUESTION NO: 6

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

Explanation:

`netadm list` is the appropriate command because it displays the network configuration profiles and their current state under Solaris network administration. It lets an administrator identify whether the system is using the Automatic network configuration profile, which is managed by NWAM, or the DefaultFixed profile, which represents manually configured network interfaces. The command's output includes profile names, profile types, and activation state, so it provides the needed operational view before configuring an Automated Installer server. For example, an active `Automatic` profile indicates that NWAM is managing the network configuration, while an active `DefaultFixed` profile indicates fixed/manual networking. This distinction is important because AI server networking commonly requires a stable, manually administered interface and address. Oracle documents `netadm` as the utility for viewing and administering active network configuration profiles, including listing their state. The related Oracle Solaris networking documentation also describes the Automatic and DefaultFixed profiles and how profile activation determines whether NWAM or fixed network configuration is in use. See the Oracle netadm(1M) manual page and the Oracle Solaris network configuration profile documentation.

QUESTION NO: 7

Which two actions permit the system-log service to receive messages from a remote Solaris host?

A. setting the property config/log_from_remote to true and restarting the service

B. setting the property config/log_from_remote to *.noticoand restart the service

C. configuring a selector for remote messages in the /etc/syslog.conf file

D. ensuring that port 514 is open to remote traffic and doesn't require a password

ANSWER: A D

Explanation:

Setting the property `config/log_from_remote` to `true` and restarting the service enables the Solaris `system-log` service to accept syslog datagrams originating on other hosts. The service is managed by SMF, and this property controls whether its syslog daemon listens for and processes remote log traffic; after changing the property, restarting (or refreshing, where appropriate) the service applies the updated configuration. The receiving system must also be reachable for the syslog transport. Ensuring that port 514 is open to remote traffic allows incoming UDP syslog packets from the remote Solaris host to reach the system-log daemon. Traditional Solaris syslog transport does not use a per-message password authentication exchange, so network reachability and the service's remote-message setting are the relevant prerequisites. Once messages are accepted, the normal `syslog.conf` rules determine how the received records are selected and written to local destinations. Oracle documents the system-log service and its remote logging controls in the system-log(8) manual page; syslog message routing and selector processing are described in the syslog.conf(5) manual page.

QUESTION NO: 8

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:

Import the service is the required next step. In Oracle Solaris SMF, the XML manifest is the service definition, but it does not become known to the Service Management Facility merely because the file and its executable method script exist on disk. The manifest must be validated and imported into the SMF repository, normally by using `svccfg import /path/to/manifest.xml`. During import, SMF reads the service and instance definitions, verifies the manifest structure and properties, and records the configuration in its repository. This establishes the service FMRI and makes it available for SMF administration.

After import completes, an administrator can use `svcadm enable` to enable the service instance. The start method specified by the manifest then invokes the executable script or command created for the service. Importing is therefore the installation/registration action that bridges the manifest file and the managed SMF service. Oracle documents `svccfg import` as the command used to import service configuration from an XML manifest into the repository; manifests may also be placed in standard manifest directories for automatic import during boot. See the svccfg(8) manual page and Oracle's SMF overview documentation.

QUESTION NO: 9

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:

`dispadmin` is the administrative utility used to configure scheduling-class dispatch parameters and to select the system-wide default scheduling class. To make the Fair Share Scheduler (FSS) the default class for newly created processes, an administrator uses `dispadmin -d FSS`. This changes the default scheduler selection so that processes which do not explicitly request another scheduling class are placed in FSS. FSS allocates CPU resources according to the shares assigned to projects, making it appropriate when CPU entitlement should be managed through project-based resource controls rather than solely through traditional time-sharing priorities.

The default-class setting is distinct from assigning an individual process to a scheduling class or modifying project attributes. It is a dispatch-configuration action performed through `dispadmin`. Oracle documents the `-d` option as specifying the default scheduling class and identifies FSS as the Fair Share Scheduler class. See the dispadmin(8) manual page and Oracle's FSS overview for the scheduler's role and configuration details.

QUESTION NO: 10

You want to configure your IPS repository server for high network bandwidth and network availability. Which two technologies are best suited for achieving these goals?

A. naxbw resource control

B. zpool disk aggregation

C. link load balance

D. link aggregation

E. IP multipathing

ANSWER: D E

Explanation:

link aggregation and **IP multipathing** are the Solaris technologies intended to improve network capacity and availability by using multiple network interfaces. A link aggregation combines several physical data links into one logical aggregation at the data-link layer. Traffic can be distributed across member links, increasing available aggregate throughput, while the aggregation can continue operating when an individual member link fails. This makes it especially appropriate where the IPS repository server must serve many clients and a suitable switch configuration is available.

IP multipathing, commonly called IPMP, groups interfaces at the IP layer to provide resilient network connectivity. IPMP detects interface or path failures and can move IP addresses to a remaining usable interface, preserving access to the server. It can also spread outbound traffic across active interfaces, improving utilization of the available network capacity. IPMP is therefore a strong choice when network availability is required and when path-level failure detection and IP-layer failover are important. Solaris supports using these technologies according to the network design: link aggregation supplies link-layer aggregation, and IPMP supplies IP-layer multipathing and failover. Oracle documents link aggregation in the Managing Network Datalinks in Oracle Solaris and IPMP in the Administering IPMP Groups.

QUESTION NO: 11

What is the effect of configuring privileges via the zonecfg utility?

A. It forces every /one process to run with the same privileges.

B. It restricts zone processes to the inherited set of zsched's privileges.

C. It restricts zone processes to the inherited set of zoneadmd's privileges.

D. It removes some privileges that are normally available in the zone.

E. It can add some new privileges to or exclude some default privileges from the zone.

ANSWER: B D E

Explanation:

Configuring a zone's privilege limit with `zonecfg` establishes the maximum privilege set that processes in that non-global zone can obtain. The zone scheduler process, `zsched`, is created for the zone with this configured limit, and processes created within the zone inherit that constrained privilege environment. Consequently, the configured limit governs the privileges available throughout the zone rather than assigning an identical active privilege set to every individual process.

The global-zone administrator can modify the zone's default privilege limit by adding permitted privileges when a workload requires capabilities not present in the default set. The administrator can also subtract privileges from the default set to enforce a tighter security boundary. Thus, privilege configuration can both remove privileges normally available to zone processes and add eligible privileges or exclude default privileges from the zone's limit set. These changes take effect as part of the zone configuration and are intended to implement least privilege while allowing explicit, controlled exceptions for zone workloads. Oracle documents the `limitpriv` zone configuration property and its role as the privilege limit set in the zonecfg(8) manual page; see also Oracle's Zones administration overview.

QUESTION NO: 12

Examine the output of the following commands:

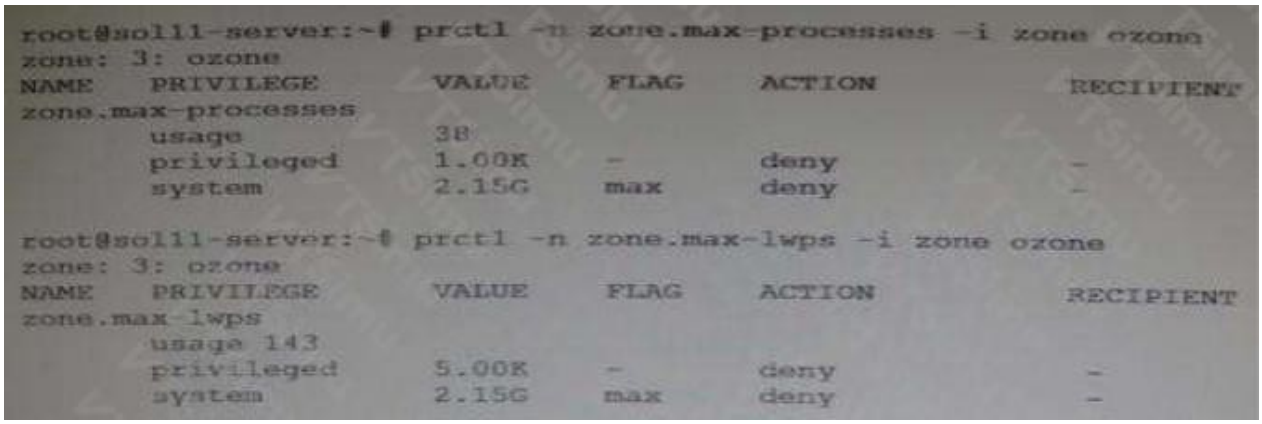

```
root@soll1-server:~# prctl -n zone.max-processes -i zone ozone
zone: 3: ozone
NAME    PRIVILEGE       VALUE    FLAG    ACTION                 RECIPIENT
zone.max-processes
        usage           38
        privileged      1.00K    -       deny                   -
        system          2.15G    max     deny                   -

root@soll1-server:~# prctl -n zone.max-lwps -i zone ozone
zone: 3: ozone
NAME    PRIVILEGE       VALUE    FLAG    ACTION                 RECIPIENT
zone.max-lwps
        usage 143
        privileged      5.00K    -       deny                   -
        system          2.15G    max     deny                   -
```

Which statement is true concerning these resource controls settings?

A. The zone.max-lwps resource control was set to its value when the zone.max-processes resource control was set to its value.

B. The prctl command can be used to change the max value for these resource controls.

C. The zone.max-processes resource control will deny zombie processes from exhausting the process table.

D. The zone.max-lwps resource control will deny zombie processes from exhausting the process table.

ANSWER: C

Explanation:

The statement “The zone.max-processes resource control will deny zombie processes from exhausting the process table.” is correct. The `zone.max-processes` resource control establishes the maximum number of process-table slots that can be simultaneously used by processes within a non-global zone. This accounting is significant because a zombie process, although it has terminated execution and no longer has an active lightweight process, still retains a process-table entry until its parent collects its termination status. Consequently, an accumulation of unreaped zombie processes consumes process-table slots just as active processes do. By enforcing the zone-level process limit, `zone.max-processes` prevents the zone from continuing to allocate process entries once its configured limit is reached, protecting the system’s process-table resources from exhaustion caused by either active or zombie processes. This control is therefore specifically suitable for constraining the process-entry impact of zombies in a zone. Oracle documents zone resource controls and their relationship to zone-wide resource consumption in the Oracle Solaris resource-controls documentation; the related administrative command behavior is described in the prctl(1) manual page.

QUESTION NO: 13

You are about to configure resource controls for a nonglobal zone. You want to first examine settings as well as the system limits for those controls. Which command fetches this information?

A. priocntl

B. zonecfg

C. rctladm

D. prctl

ANSWER: B

Explanation:

`zonecfg`, using its `info` subcommand, is the appropriate command when preparing to configure resource controls for a nonglobal zone. The zone configuration is maintained and administered through `zonecfg`; its `info` operation displays the current configuration, either in full or for a specified resource type. This lets an administrator inspect configured zone-wide controls, such as caps and limits, before making changes. For example, after selecting a zone with `zonecfg -z zonename`, an administrator can use `info` to review the zone definition, or use `info capped-memory` or another relevant resource type to focus on a particular configured resource. Reviewing this information first is important because resource controls are applied as part of the zone's persistent configuration and take effect when the zone is booted or rebooted as required by the specific property. Oracle documents `zonecfg` as the utility for configuring zones and describes `info` as the command that displays current configuration information. See the zonecfg(1M) manual page and Oracle's Zones configuration overview.

QUESTION NO: 14

You have assigned rights profiles directly to the uses frank and now you want to add another profile. Which command enables you to list profiles directly assigned to frank?

A. userattr profiles frank

B. profiles –p frank

C. useratter –p frank

D. profiles frank

ANSWER: A

Explanation:

`userattr profiles frank` is correct because `userattr` retrieves a named user attribute from the Solaris user-attribute databases for the specified account. The `profiles` attribute contains the rights profiles that are assigned directly to the user, so this command displays the profile list currently recorded for `frank`. This is the appropriate information to inspect before modifying the account's directly assigned rights profiles, such as with the `usermod -P` command.

Solaris RBAC stores user-specific authorizations, profiles, roles, and related attributes in user attribute data, commonly maintained through the local `user_attr` database or a configured naming service. Querying the specific `profiles` attribute is therefore more precise than requesting general profile information: it identifies the direct assignments on the user account itself. The command syntax is `userattr attribute_name user`, making `profiles` the attribute name and `frank` the target user. Oracle documents this syntax and the purpose of the `profiles` attribute in the userattr(1) manual page and its Solaris RBAC administration guidance in Role-Based Access Control.

QUESTION NO: 15

Link aggregation and IP multipathing both offer some benefits for network performance and reliability. Identify two correct statements.

A. IPMP requires full duplex, point-to-point links.

B. A router is a single point of failure for link aggregation.

C. Link aggregation allows a standby interface to be automatically enabled if another interface fails.

D. Depending on load balancing algorithms, packets may not be balanced among all IPMP active interfaces.

E. Link aggregation uses additional interfaces to improve performance, without requiring additional IP address.

ANSWER: C E

Explanation:

Link aggregation allows a standby interface to be automatically enabled if another interface fails. In an LACP-managed aggregation, the aggregation can have more eligible links than can be active simultaneously. Links that cannot currently participate as active members can be maintained in standby status. If an active member fails, a suitable standby link can be selected and activated, preserving connectivity and helping maintain the aggregate's available bandwidth.

Link aggregation uses additional interfaces to improve performance, without requiring additional IP address. The aggregation presents its grouped physical network interfaces as one logical data link at the MAC layer. An IP interface and its address are configured on that logical aggregation rather than separately on every member port. Consequently, traffic can use the capacity of multiple physical links while the host retains a single IP-level attachment for that aggregation. This also provides link-level resiliency because traffic can continue through remaining active member links when a member becomes unavailable. Oracle's Solaris networking documentation describes link aggregation as combining multiple interfaces into a single logical link for increased throughput and availability, while distinguishing it from IP multipathing, which operates at the IP layer. See Oracle's Link Aggregations overview and IPMP administration documentation.

QUESTION NO: 16

You configured a limit of 100 LWPs project. You want to ensure that the LWP limit was not set too low, so you need to monitor the LWPs currently in use by the project.

Which two options could you use to monitor the current LWP resource control and the consumption of resources for this project?

A. prtcl $$

B. configuring syslogd to log messages received from the resource manager daemon

C. ps –o taskid –p

D. prctl –n task.max-lwps $$

E. rctladm –l task.max-lwps

F. rctladm –e syslog task.max-lwps; when the threshold for the resource is exceeded, a log entry will be generated by syslogd

ANSWER: D F

Explanation:

`prctl -n task.max-lwps $$` displays the `task.max-lwps` resource control associated with the task containing the current shell. Its output includes the configured control values and a `usage` value, which is the current number of lightweight processes consumed by that task. Running the command from a process in the relevant project therefore provides a direct way to observe present LWP consumption against the configured project task limit. The `$$` shell parameter supplies the PID of the current shell, allowing `prctl` to identify its task and applicable resource controls.

`rctladm -e syslog task.max-lwps` enables the system-wide `syslog` action for this resource control. When a task reaches or attempts to exceed the applicable `task.max-lwps` threshold, Solaris generates a syslog event. These events provide ongoing operational monitoring and an audit trail of limit-pressure or denial events, helping an administrator determine whether the configured maximum of 100 LWPs is being encountered in real workloads. The resource-control action is enabled globally, while the control is evaluated for each affected task. Oracle documents `prctl` resource-control reporting in the prctl(1) manual page and global resource-control actions in the rctladm(8) manual page.

QUESTION NO: 17

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: C D

Explanation:

`/etc/security/prof_attr` is the rights-profile attribute database. A profile entry in this database defines the profile name and its attributes, including the `auths` attribute used to associate Solaris authorizations with that rights profile. This is therefore the appropriate profile database for adding the selected authorizations.

`/etc/security/exec_attr` is the execution-attribute database. It maps commands assigned through a rights profile to execution security attributes. The `privs` execution attribute can specify the additional privileges with which a command is run, making this the appropriate database for commands in the profile that require extra privileges.

On Solaris 11, these named local databases are assembled from fragment files. Administrative additions are normally placed in the applicable `local-entries` fragment directories, such as `/etc/security/prof_attr.d/local-entries` and `/etc/security/exec_attr.d/local-entries`, rather than editing package-delivered content. The profile and execution-attribute databases nonetheless identify the two RBAC data sources that hold the required authorization and privilege definitions. Oracle documents the profile attributes in prof_attr(5) and command execution attributes, including privilege settings, in exec_attr(5).

QUESTION NO: 18

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 are configured for link-based detection by default.

D. You must first disable link based detection before configuring probe-based failure detection.

ANSWER: C

Explanation:

GLDv3 NICs are configured for link-based failure detection by default. In Oracle Solaris IPMP, probe-based failure detection and link-based failure detection are complementary mechanisms managed by the `in.mpathd` daemon. Probe-based detection verifies end-to-end reachability by sending ICMP probes, either using explicitly configured test addresses or through transitive probing when no test addresses are present. Separately, when an underlying interface and its NIC driver support reporting link state, IPMP automatically uses link-state notifications to recognize physical link failures promptly. Thus, enabling probe-based detection does not require disabling link-based detection; both can be active for the same IPMP group. GLDv3 is the current Solaris Generic LAN Driver framework for Ethernet drivers and supplies the capabilities needed for modern NIC drivers to report link-state changes to the networking stack. This automatic link-state capability improves failure recognition for cable, switch-port, and similar physical connectivity events, while probing covers reachability failures that may occur even when a physical link remains up. Oracle documents IPMP failure detection and the relationship between probe-based and link-based mechanisms in its Solaris networking documentation. See the Oracle Solaris 11.4 Network Administration documentation and the in.mpathd(8) manual page.