

1Z0-821^{Q&As}

Oracle Solaris 11 System Administration

Pass Oracle 1Z0-821 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass2lead.com/1z0-821.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



QUESTION 1

Identify three options that describe the new Oracle Solaris 11 zone features.

- A. There are boot environments for zones.
- B. Administrators can delegate common administration tasks by using RBAC.
- C. Oracle Solaris 11 supports Solaris 8, 9, and 10 branded zones.
- D. You can migrate a physical Solaris 10 system and its non-global zones to a solaris10 branded zone running on an Oracle Solaris 11 system.
- E. It is possible to change the host ID of a zone.

Correct Answer: ABD

A: The beadm utility includes support for creating and administering non-global zone boot environments. Note: A boot environment is a bootable instance of the Oracle Solaris operating system image plus any other application software packages installed into that image. System administrators can maintain multiple boot environments on their systems, and each boot environment can have different software versions installed.

B: Role-based access control (RBAC) is a security feature for controlling user access to tasks that would normally be restricted to the root role. By applying security attributes to processes and to users, RBAC can divide up superuser capabilities among several administrators.

QUESTION 2

You want the system to generate an email notification each time one of the services has changed its state. Which option would send an email message to the system administrator whenever a service changes to the maintenance state?

- A. Use the setsc command in ALOM to enable the mail alerts to be sent to a specified email address whenever the fault management facility detects a service change to the maintenance state.
- B. Make an entry in the /etc/syslog.conf file to instruct syslogd to send an email alert when it receives a message from the SMF facility that a service has changed to the maintenance state.
- C. Use the svccfg setnotify command to create a notification and send an email when a service enters the maintenance state.
- D. Use the scvadm command to enable the notification service. Set the ? maintenance option on the netnotify service to send an email when a service enters the maintenance state.

Correct Answer: C

This procedure causes the system to generate an email notification each time one of the services or a selected service has a change in state. You can choose to use either SMTP or SNMP. Normally, you would only select SNMP if you already have SNMP configured for some other reason.

By default, SNMP traps are sent on maintenance transitions. If you use SNMP for monitoring, you can configure additional traps for other state transitions.

1. Become an administrator or assume a role that includes the Service Management rights profile.

2. Set notification parameters.

Example: The following command creates a notification that sends email when transactions go into the maintenance state # /usr/sbin/svccfg setnotify -g maintenance mailto:sysadmins@example.com

QUESTION 3

When you issue the "gzip: zommand not found" message is displayed. You need to install the gzip utility on your system.

Which command would you use to check if the gzip utility is available from the default publisher for installation?

- A. pkg info|grep gzip
- B. pkg list SUNWgzip
- C. pkg contents gzip
- D. pkg search gzip

Correct Answer: D

Searching for Packages

Use the pkg search command to search for packages whose data matches the specified pattern.

Like the pkg contents command, the pkg search command examines the contents of packages. While the pkg contents command returns the contents, the pkg search command returns the names of packages that match the query.

pkg search search [-Hlaf|pr] [-o attribute ...] [-s repo_uri] query Search for matches to the query, and display the results. Which tokens are indexed are action-dependent, but may include content hashes and pathnames.

Note: pkg is the retrieval client for the image packaging system. With a valid configuration, pkg can be invoked to create locations for packages to be installed, called \\images\\, and install packages into those images. Packages are published by publishers, who may make their packages available at one or more repositories. pkg, then, retrieves packages from a publisher\\'s repository and installs them into an image.

QUESTION 4

Consider the following commands:

```
rm file1
echo "Hello, world" > file2
cat file1 && cat file2
```

What is displayed when this sequence of commands is executed using the bash shell?

- A. Hello, world
- B. cat: cannot open file1

- C. cat: cannot open file1Hello, world
- D. cat: cannot open file1 Hello, World
- E. bash: syntax error near unexpected token `\'andand\'

Correct Answer: B

First line (rm file1) deletes/removes file1.

Second line captures the text into file2.

The first part of line 3 (cat file1) fails as the file1 does not exist. The andand (AND) operator will ensure that the third line fails. The result of line 3 will be the result of first part of line 3 (cat file1).

Note: cat - concatenate files and print on the standard output

Note #1: A list is a sequence of one or more pipelines separated by one of the operators `;\', `and\', `andand\', or `||\', and optionally terminated by one of `;\', `and\', or a newline.

Of these list operators, `andand\' and `||\' have equal precedence, followed by `;\' and `and\', which have equal precedence.

AND and OR lists are sequences of one or more pipelines separated by the control operators `andand\' and `||\', respectively. AND and OR lists are executed with left associativity.

An AND list has the form

command1 andand command2

command2 is executed if, and only if, command1 returns an exit status of zero.

An OR list has the form

command1 || command2

command2 is executed if, and only if, command1 returns a non-zero exit status.

The return status of AND and OR lists is the exit status of the last command executed in the list.

Note #2 (on exit status): Zero means command executed successfully, if exit status returns non-zero value then your command failed to execute.

QUESTION 5

You are logged in to a Solaris 11 system as user jack. You issue the following sequence of commands: Identify two correct statements.

```
jack@solaris:~$ id
uid=65432 (jack) gid=10(staff) groups=10(staff)
jack@solaris:~$ su
Password:
jack@solaris:~#
```

- A. You have the effective privilege of the account root.
- B. Your GID is 10.
- C. Your home directory is /root.
- D. You are running the shell specified for the account root.
- E. Your UID is 1.

Correct Answer: AB

Oracle Solaris provides predefined rights profiles. These profiles, listed in the /etc/security/prof_attr, can be assigned by the root role to any account. The root role is assigned all privileges and all authorizations, so can perform all tasks, just as root can when root is a user.

To perform administrative functions, you open a terminal and switch the user to root. In that terminal, you can then perform all administrative functions.

\$ su - root Password: Type root password # When you exit the shell, root capabilities are no longer in effect.

QUESTION 6

You have a process called bigscript, and you need to know the PID number for this process.

Which command will provide that information?

- A. pkill bigscript
- B. ps bigscript
- C. pgrep bigscript
- D. prstat bigscript

Correct Answer: C

Pgrep takes a process name and return a PID.

Note: pgrep looks through the currently running processes and lists the process IDs which matches the selection criteria to stdout. All the criteria have to match. For example, pgrep - u root sshd will only list the processes called sshd AND owned by root.

Incorrec answers:

ps bigscript: You can't pass a name to ps, it interprets it as arguments.

QUESTION 7

You have a user that needs to use the cron tool to schedule some repetitive tasks. When the user enters the crontab ? command in a terminal window, the following error appears:

crontab: you are not authorized to use cron. Sorry

In order to troubleshoot this issue, in what directory would you start your invest

- A. /etc/cron.d
- B. /var/spool/cron
- C. /var/spool/cron/crontable
- D. /var/spool/cron/atjobs

Correct Answer: A

crontab: you are not authorized to use cron. Sorry.

This message means that either the user is not listed in the cron.allow file (if the file exists), or the user is listed in the cron.deny file.

You can control access to the crontab command by using two files in the /etc/cron.d directory: cron.deny and cron.allow. These files permit only specified users to perform crontab command tasks such as creating, editing, displaying, or removing their own crontab files.

The cron.deny and cron.allow files consist of a list of user names, one user name per line.

QUESTION 8

Which two statements are true concerning the network stack on Oracle Solaris 11?

- A. Hardware network interfaces and datalinks have a one-to-one relationship.
- B. IP addresses are assigned to datalinks.
- C. A single IP interface can have either an IPv4 address or an IPv6 address but not both.
- D. A single IP interface can have both an IPv4 address and an IPv6 address.
- E. A single datalink can have only one IP interface.

Correct Answer: AD

QUESTION 9

Your SPARC server will not boot into multi user-server milestones and you need to troubleshoot to out why. You need to start the server with minimal services running so that you can go through each milestone manually to troubleshoot the issue.

Select the option that boots the server with the fewest services running.

- A. boot -s
- B. boot milestone none
- C. boot -m milestone=single-user
- D. boot -m milestone=none
- E. boot -m none

Correct Answer: D

The command boot -m milestone=none is useful in repairing a system that have problems booting early.

Boot Troubleshooting: To step through the SMF portion of the boot process, start with: boot -m milestone=none Then step through the milestones for the different boot levels: svcadm milestone svc:/milestone/single-user:default svcadm milestone svc:/milestone/multi-user:default svcadm milestone svc:/milestone/multi-user-server:default

QUESTION 10

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

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

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

Correct Answer: CE

In this output, the statistics for interrupt (INTRS) are significant. Low interrupt numbers indicate greater efficiency in performance. If the interrupt numbers are high, then you might need to add more resources to the specific link.

Example:

```
# dlstat -r -i 1
```

```
LINK IPKTS RBYTES INTRS POLLS CH50
```

```
e1000g0 101.91K 32.86M 87.56K 14.35K 3.70K 205 5
```

nxge1 9.61M 14.47G 5.79M 3.82M 379.98K 85.66K 1.64K

vnic1 8 336 0 0 0 0 0

e1000g0 0 0 0 0 0 0 0

nxge1 82.13K 123.69M 50.00K 32.13K 3.17K 724 24

vnic1 0 0 0 0 0 0 0

Note: `dlstat show-link [-r [-F] | -t] [-i interval] [-a] [-p] [-o field[, ...]] [-u R|K|M|G|T|P] [link]` Display statistics for a link.

`-r`

Display receive-side statistics only. Includes bytes and packets received, hardware and software drops, and so forth.

List of supported RX fields:

link

iusedby

ibytes

ipkts

intrs

polls hdrops: hardware drops sdrops: software drops (owing to bandwidth enforcement) ch50: number of packet chains of length > 50

QUESTION 11

The ZFS configuration on your server is:

Pool1 6.67G31K/pool

Pool1/data31K31K/data

Select the three commands that you would use to 1. Create, 2. List, and 3. Delete a snapshot of the /data file system.

- A. `zfs snapshot pool1/data@now`
- B. `zfs create snapshot pool1/data@now`
- C. `zfs list -t snapshot`
- D. `zfs list -t snapshot pool1/data`
- E. `zfs destroy pool1/data@now`

F. zfs destroy snapshot pool1/data@now

Correct Answer: ADE

A: Snapshots are created by using the zfs snapshot command, which takes as its only argument the name of the snapshot to create.

D: You can list snapshots as follows: # zfs list -t snapshot

E: Snapshots are destroyed by using the zfs destroy command. For example:

```
# zfs destroy tank/home/ahrens@now
```

QUESTION 12

Given:

file1 and file2 are text files.

dir1 and dir2 are directories.

Which two commands will be successful?

A. cp dir1 dir1

B. cp dir1 file1

C. cp file? dir1

D. cp file. dir1

E. cp file% dir2

F. cp file1 file2 dir1

Correct Answer: CF

C: Here the wildcard character ? is used (Matches any single character).

file1 and file2 will be copied into dir1

F: the two files file1 and file2 are copied into directory dir1.

Note: cp - copy files and directories

Copy SOURCE to DEST, or multiple SOURCE(s) to DIRECTORY.

Cp has three principal modes of operation. These modes are inferred from the type and count of arguments presented to the program upon invocation.

*

When the program has two arguments of path names to files, the program copies the contents of the first file to the second file, creating the second file if necessary.

*

When the program has one or more arguments of path names of files and following those an argument of a path to a directory, then the program copies each source file to the destination directory, creating any files not already existing.

*

When the program's arguments are the path names to two directories, cp copies all files in the source directory to the destination directory, creating any files or directories needed. This mode of operation requires an additional option flag, typically r, to indicate the recursive copying of directories. If the destination directory already exists, the source is copied into the destination, while a new directory is created if the destination does not exist.

QUESTION 13

You wish to edit your crontab file that is located in /var/spool/cron/crontab. What command must you enter to edit this file?

- A. crontab -e
- B. crontab -e /var/spool/cron/crontab
- C. crontab -r
- D. crontab -e /etc/default/cron

Correct Answer: A

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

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

QUESTION 14

Which files must be edited in order to set up logging of all failed login attempts?

- A. /etc/default/login, /var/adm/loginlog, /etc/syslog.conf
- B. /etc/default/login, /var/adm/authlog, /etc/syslog.conf
- C. /var/adm/loginlog, /var/adm/authlog, /etc/syslog.conf
- D. /etc/default/login, /var/adm/authlog, /var/adm/loginlog

Correct Answer: B

This procedure captures in a syslog file all failed login attempts.

1.

Set up the /etc/default/login file with the desired values for SYSLOG and SYSLOG_FAILED_LOGINS Edit the

/etc/default/login file to change the entry. Make sure that SYSLOG=YES is uncommented.

2.

Create a file with the correct permissions to hold the logging information. Create the authlog file in the /var/adm directory.

3.

Edit the syslog.conf file to log failed password attempts. Send the failures to the authlog file.

QUESTION 15

The core dump configuration in your non global zone is

```
global core file pattern: /var/core/core.%f.%p
global core file content: default
init core file pattern: /var/core/pprocess/core.%f.%p
init core file content: default
global core dumps: enabled
per-process core dumps: enabled
global setid core dumps: disabled
per-process setid core dumps: disabled
global core dump logging: disabled
```

A user is running a process in a non-global zone (testzone) and the process crashes. The process information is:

user126632618017:46:42pts/20:00/usr/bin/bash

When the user's process crashes in testzone, a non-global zone, where will the core dump be saved?

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

Correct Answer: E

The line init core file pattern: /var/core/core.%f.%p will be used for the non-global process to determine the destination of the dump file.

Note: When a process is dumping core, up to three core files can be produced: one in the per-process location, one in the system-wide global location, and, if the process was running in a local (non-global)

zone, one in the global location for the zone in which that process was running.