Oracle 1Z0-821 Mock Exam

Shared by Watkins on 17-06-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You created an IP address for interface not.3 with the following command, which executed successfully:

ipadm create-addr --T static --a 192.168.0.100/24 net3/v4

You then ran:

ipadm show--if

The result indicated that the interface was down.

You then ran:

ipadm delete-addr net3/v4

ipadm create-addr --T static --a 192.168.0.101/24 net3/v4

ipadm show-if

The last command indicated that the interface was up.

Why did it work with the second address specified, but not the first?

Options:

A- The 192.168.0.100 address is reserved for broadcast messages.
B- Another device exists on the network, using the 192.168.0.100 address.
C- The network interface card does not support the address 192.168.0.100.
D- The address 192.168.0.100 is at a boundary and may not be configured in Oracle Solaris 11.
E- 192.168.0.100 is a DHCP address and may not be statically configured in Oracle Solaris 11.

Answer:

B

Explanation:

The first IP address is already in use.

Question 2

Question Type: MultipleChoice

The core dump configuration for your system is:

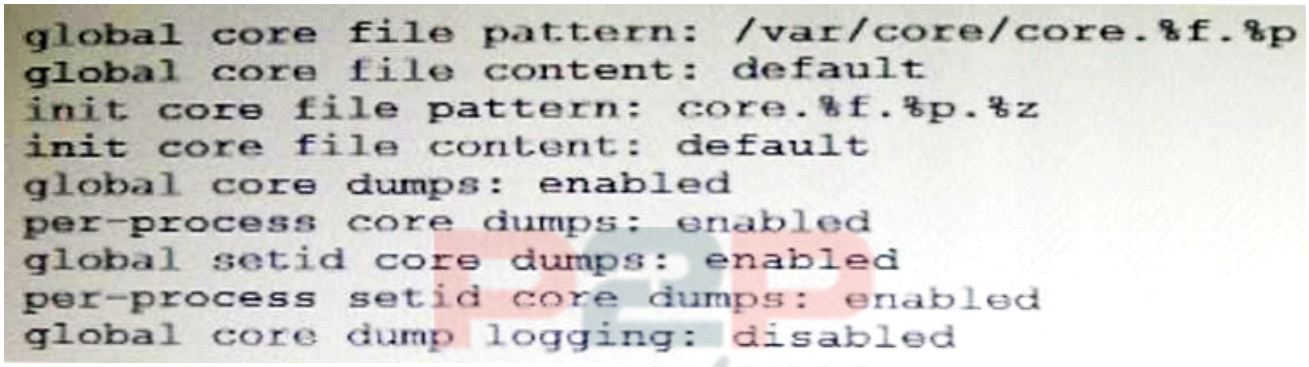

A user is running a process in the global zone and the process crashes. The process information is:

User1 2663 2618 0 17:46:42 pts/2 0:00 /usr/bin/bash

The server host name is: zeus

What will the per-process core file be named?

Options:

A- core.bash.2663.global
B- core.bash.2663.zeus
C- /var/core/core.bash.2663
D- /var/core/core.bash.2663.global

Answer:

C

Explanation:

Note the first line:

global core file pattern: /globalcore/core.%f.%p

The program name is bash

The runtime process ID is 2663

Note: By default, the global core dump is disabled. You need to use thecoreadmcommand with

the-e globaloption to enable it. The-goption causes the command to append the program name(%f)and the runtime process ID(%p)to the core file name.

Question 3

Question Type: MultipleChoice

You need to configure an ISCSI target device on your x86 based Oracle Solaris II system. While configuring the iSCSI device, the following error is displayed:

bash: stmfadm: command not found

Which option describes the solution to the problem?

Options:

A- The COMSTAR feature is not supported on the x86 platform. The feature is supported only on the SPARC platform.
B- Use the iscsitadm command on the x86 platform when configuring an iSCSI target.
C- Install the storage-server group package on this system.
D- Start the iSCSI target daemon on this system.

Answer:

C

Explanation:

STMF -- Manages transactions, such as context and resources for Small Computer System Interface (SCSI) command execution, and tracking logical unit and port providers. STMF also handles logical unit mappings, allocating memory, recovering failed operations, enumeration, and other necessary functions of an I/O stack.

STMF is controlled by stmfadm, and stmfadm is the majority of the commands you will be using to administer COMSTAR (COmmon Multiprotocl Scsi TARget).

Install the packages you need for COMSTAR with iSCSI and reboot:

pfexec pkg install storage-server

pfexec pkg install SUNWiscsit

shutdown -y -i6 -g0

Note: You can set up and configure a COMSTAR Internet SCSI (iSCSI) target and make it available over the network. The iSCSI features can work over a normal Internet connection (such as Ethernet) using the standard iSCSI protocol. The iSCSI protocol also provides naming and discovery services, authentication services using CHAP and RADIUS, and centralized management through iSNS.

The COMSTAR target mode framework runs as the stmf service. By default, the service is disabled. You must enable the service to use COMSTAR functionality. You can identify the service with the svcs command. If you have not rebooted the server since installing the group/feature/storage-server package, the service might not be enabled correctly.

Question 4

Question Type: MultipleChoice

User1 is attempting to run the following command:

cp bigfile verybig

The system displays the following errer:

cp: cannot create verybig: Disc quota exceeded

Your initial troubleshooting shows that the df -h command indicates the account is at 100% capacity. What command would you use to determine how much disk space the user has available?

Options:

A- zfs get quota rpool/export/home/user1
B- zfs userused@user1
C- zfs quota=1M /rpool/export/home/user1
D- df --h | grep user1

Answer:

A

Explanation:

ZFS quotas can be set and displayed by using the zfs set and zfs get commands. In the following example, a quota of 10 Gbytes is set on tank/home/bonwick.

zfs set quota=10G tank/home/bonwick

zfs get quota tank/home/bonwick

NAME PROPERTY VALUE SOURCE

tank/home/bonwick quota 10.0G local

Question 5

Question Type: MultipleChoice

The storage pool configuration on your server is:

```
pool1           200K        3.91G    31K    /pool1
pool1/data       31K        3.91G    31K    /pool1/data
pool1           ONLINE         0      0      0
   c4t0d0       ONLINE         0      0      0
```

You back up the /pool1/data file system, creating a snapshot and copying that snapshot to tape (/dev/rmt/0). You perform a full backup on Sunday night and Incremental backups on Monday through Saturday night at 11:00 pm. Each incremental backup will copy only the data that has been modified since the Sunday backup was started.

On Thursday, at 10:00 am, you had a disk failure. You replaced the disk drive (c4t0d0). You created pool (pool1) on that disk.

Which option would you select to restore the data in the /pool1/data file system?

Options:

A- zfs create pool1/dataLoad the Monday tape and enter:zfs recv pool1/data < /dev/rmt/0Load the Wednesday tape and enter:zfs recv --F pool1/data < /dev/rmt/0
B- Load the Sunday tape and restore the Sunday snapshot:zfs recv pooll/data < /dev/rmt/0zfs rollback pool1/data@monLoad the Wednesday tape and restore the Wednesday snapshot:zfs recv --i pooll/data < /dev/rmt/0zfs rollback pool1/data@wed
C- zfs create pooll/dataLoad the Wednesday tape and enter:zfs recv -F pool1/data < /dev/rmt/0
D- Load the Sunday tape and enter:zfs recv pool1/data < /dev/rmt/0Load the Wednesday tape and enter:* commands missing*

Answer:

D

Explanation:

First the full backup must be restored. This would be the Sunday backup.

Then the last incremental backup must be restored. This would be the Wednesday backup.

Before restoring the Wednesday incremental file system snapshot, the most recent snapshot must first be rolled back.

By exclusion D) would be best answer even though it is incomplete.

Question 6

Question Type: MultipleChoice

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?

Options:

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 Deduplication Property

Solaris Express Community Edition, build 129: In this Solaris release, you can use the deduplication property to remove redundant data from your ZFS file systems. If a file system has the dedup property enabled, duplicate data blocks are removed synchronously. The result is that only unique data is stored and common components are shared between files.

You can enable this property as follows:

```
# zfs set dedup=on tank/home
```

P2P exams

P2P exams

To Get Premium Files for 1Z0-821 Visit
https://www.p2pexams.com/products/1z0-821

For More Free Questions Visit
https://www.p2pexams.com/oracle/pdf/1z0-821