



1Z0-820^{Q&As}

Upgrade to Oracle Solaris 11 System Administrator

Pass Oracle 1Z0-820 Exam with 100% Guarantee

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

<https://www.passapply.com/1z0-820.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

You are troubleshooting network throughput on your server.

To confirm that the load balancing among aggregated links is functioning properly, you want to examine the traffic statistics on the links comprising the aggregation.

The correct command is _____.

- A. `dlstat - aggr`
- B. `dlstat show-aggr`
- C. `dlstat show-link -r`
- D. `dlstat show-link -aggr`
- E. `dlstat show-phys -aggr`

Correct Answer: B

Explanation: `dlstat show-aggr [-r | -t] [-i interval] [-p] [-o field[,...]] [-u R|K|M|G|T|P] [link]` Display per-port statistics for an aggregation.

QUESTION 2

You are troubleshooting the failure of a computer to mount an NFS file system hosted by a server (hostname mars) in the local area network.

Select the three commands that will enable you to identify the problem.

- A. `ping -s mars`
- B. `cat /etc/vfstab`
- C. `cat /etc/dfs/dfstab`
- D. `sharemgr show -v`
- E. `showmount -e mars`
- F. `rpcinfo -s mars | egrep `nfs|mountd\``

Correct Answer: AEF

mount: mount-point /DS9 does not exist.

To solve the mount point error condition, check that the mount point exists on the client. Check the spelling of the mount point on the command line or in the `/etc/vfstab` file (B) on the client, or comment out the entry and reboot the system.



Note: The `/etc/vfstab` file lists all the file systems to be automatically mounted at system boot time, with the exception of the `/etc/mnttab` and `/var/run` file systems.

E: showmount This command displays all clients that have remotely mounted file systems that are shared from an NFS server, or only the file systems that are mounted by clients, or the shared file systems with the client access information. The command syntax is:

```
showmount [ -ade ] [ hostname ]
```

where `-a` prints a list of all the remote mounts (each entry includes the client name and the directory), `-d` prints a list of the directories that are remotely mounted by clients, `-e` prints a list of the files shared (or exported), and `hostname` selects the NFS server to gather the information from. If `hostname` is not specified the local host is queried.

F: * mountd Daemon

This daemon handles file-system mount requests from remote systems and provides access control. The

`mountd` daemon checks `/etc/dfs/sharetab` to determine which file systems are available for remote

mounting and which systems are allowed to do the remote mounting.

* Commands for Troubleshooting NFS Problems

These commands can be useful when troubleshooting NFS problems.

`rpcinfo` Command

This command generates information about the RPC service that is running on a system.

Reference: `man showmount`

QUESTION 3

Consider the following rule file for use with the Basic Audit Reporting Tool (BART). CHECK all IGNORE `dirmtime` `/etc/security` `/etc/notices` IGNORE contents `/export/home` IGNORE `mtime` size contents `/var` CHECK You are using BART to detect inappropriate changes to the file system. Identify the two correct statements describing the attributes recorded.

- A. `/var/dhcp` Attribute: size uid gid mode acl
- B. `/etc/hosts` Attributes: size uid gid mode acl intime dest
- C. `/var/spool/mqueue` Attribute: size uid gid mode acl dirmtime
- D. `/etc/security/exec_attr` Attribute: size uid mode acl mtime devnode
- E. `/export/home/kate/.profile` Attributes: uid gid mode acl dirmtime
- F. `/export/home/rick/.profile` Attributes: size uid gid mode acl mtime contents

Correct Answer: DF

Explanation: D: According to line `/etc/security`



F: According to line /export/home Not E: According to line IGNORE dirmtime Note: In default mode, the bart compare command, as shown in the following example, checks all the files installed on the system, with the exception of modified directory timestamps (dirmtime):

CHECK all IGNORE dirmtime

Note 2: The Basic Audit Reporting Tool (BART) feature of Oracle Solaris enables you to comprehensively validate systems by performing file-level checks of a system over time. By creating BART manifests, you can easily and reliably gather information about the components of the software stack that is installed on deployed systems.

BART is a useful tool for integrity management on one system or on a network of systems.

Reference: Oracle Solaris Administration: Security Services, BART Manifests, Rules Files, and Reports (Reference)

QUESTION 4

When you issue the gzip command, the "gzip: command 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

QUESTION 5

You are going to use the Automated installer (AI) to install a non global zone named zone1. You have created a custom manifest for the non-global zone and named it zone1manifest

Which command will you use to add this custom manifest to the s11-sparc install service and associate this custom manifest with the non-global zone?

- A. installadm create-profile -n s11-sparc -f /tmp/zone1manifest.xml -c
- B. installadm create-manifest -n s11-sparc -f /tmp/zone1manifest.xml -m
- C. installadm create-client -n s11-sparc -f /tmp/zone1manifest.xml -m zone1manifest -c zonename="zone1"
- D. installadm create-service -n s11-sparc -f /tmp/zone1manifest.xml -m zone1manifest -c zonename="zone1"

Correct Answer: B

Explanation: installadm add-manifest

Associates manifests with a specific install service, thus making the manifests available on the network, independently from creating a service. When publishing a non-default manifest, it is required to associate criteria either via criteria entered on the command line (-c) or via a criteria XML file (-C). Reference: installadm



QUESTION 6

You are asked to troubleshoot networking issues on an unfamiliar system.

Select the correct command to display what network devices are installed.

- A. ifconfig -a
- B. dladm show-dev
- C. dladm show-phys
- D. dladm show-ether
- E. netadm show-dev
- F. netadm show-ether

Correct Answer: B

Explanation: dladm show-dev

Shows information for all devices or the specified device.

Reference: man dladm

QUESTION 7

Your task is to convert a JumpStart sysidcfg file to an Automated Installer (AI) sc_profile.xml file, using

js2ai. Select two unsupported items that will require changes.

(Choose two.)

- A. terminal=xterms
- B. name_service=NIS+
- C. timezone=US/Pacific
- D. system_locale=en_US
- E. network_interface=PRIMARY
- F. root_password=rJmv5LUXM10cU

Correct Answer: AD

QUESTION 8

The current ZFS configuration on your server is:



pool1200K3.91G31K/pool1 pool1/data31K3.91G31K/data

Your backup policy states that you are to perform a full backup of /data on Sunday and incremental backups on Monday through Saturday. Each incremental will back up only the data that has been created or modified since the Sunday backup was started. The file systems must be available to users at all times and the data will be backed up to tape.

Which option should you choose for the Wednesday backup?

- A. On Sunday `zfs snapshot pool1/data@sunday` `zfs send pool1/data@sunday > /dev/rmt/0` On Wednesday `zfs snapshot pool1/data@wednesday` `zfs send -i pool1/data@sunday pool1/data@wednesday > /dev/rmt/0`
- B. On Sunday `zfs snapshot pool1/data@sunday > /dev/rmt/0` On Wednesday `zfs snapshot -i pool1/data@wednesday >/dev/rmt/0`
- C. On Sunday `zfs snapshot pool1/data@sunday` `zfs send pool1/data@sunday > /dev/rmt/0` On Wednesday `zfs snapshot pool1/data@wednesday` `zfs send -i pool1/data@wednesday pool1/data@sunday > /dev/rmt/0`
- D. On Sunday `zfs snapshot pool1/data@sunday` `zfs send pool1/data@sunday > /dev/rmt/0` On Wednesday `zfs snapshot pool1/data@wednesday` `zfs send -i pool1/data@wednesday > /dev/rmt/0`
- E. On Sunday `zfs snapshot pool1/data@sunday` `zfs send pool1/data@sunday | zfs recv > /dev/rmt/0` On Wednesday `zfs snapshot pool1/data@wednesday` `zfs send -i pool1/data@sunday pool1/data@wednesday|zfs recv > /dev/rmt/0`

Correct Answer: C

QUESTION 9

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



```
pool: pool1
state: ONLINE
scan: none requested
config:
```

NAME	STATE	READ	WRITE	CKSUM
pool1	ONLINE	0	0	0
mirror-0	ONLINE	0	0	0
c4t0d0	ONLINE	-	-	-
c4t1d0	ONLINE	-	-	-
mirror-1	ONLINE	0	0	0
c4t2d0	ONLINE	-	-	-
c4t3d0	ONLINE	-	-	-

errors: No known data errors

```
pool: rpool
```

```
pool: rpool
state: ONLINE
scan: none requested
config:
```

NAME	STATE	READ	WRITE	CKSUM
rpool	ONLINE	0	0	0
c3t0d0s0	ONLINE	0	0	0

errors: No known data errors

NAME	USED	AVAIL	REFER	MOUNTPOINT
pool1	138K	7.81G	32K	/pool1
pool1/prod_data	31K	7.81G	31K	/prod_data
rpool	11.6G	4.04G	34.5K	/rpool
rpool/ROOT	9.94G	4.04G	31K	legacy
rpool/ROOT/solaris	9.94G	4.04G	9.70G	/
rpool/dump	630M	4.05G	611M	-
rpool/export	6.07M	4.04G	32K	/export
rpool/export/home	6.04M	4.04G	32K	/export/home



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 pool2 zfs set mountpoint=/dev_data pool2/prod_data
- B. zpool detach pool1 pool2 zpool attach pool2 zfs set mountpoint=/dev_data pool2/prod_data
- C. zfs 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

Correct Answer: A

QUESTION 10

The COMSTAR framework provides support for the iSCSI protocol.

Identify three true statements describing the COMSTAR framework. (Choose three.)

- A. One IP port can handle multiple iSCSI target devices.
- B. The disk volume provided by the server is referred to as the LUN and can be accessed by an iSCSI initiator.
- C. The disk volume provided by the server is referred to as the target.
- D. The iSCSI initiator is configured on the server. E The iSCSI initiator is configured on the client.

Correct Answer: ACD

QUESTION 11

The COMSTAR framework provides support for the iSCSI protocol.

Select three options that correctly describe the COMSTAR framework.

- A. iSCSI devices can be used as dump devices.
- B. SCSI commands are carried over IP networks and enable you to mount disk devices from across the network onto your local system.
- C. Large amounts of data can be transferred over an IP network with very little network degradation.
- D. COMSTAR allows you to convert any Solaris11 host into a SCSI target device that can be accessed over a storage network.
- E. One IP port can handle multiple iSCSI target devices.

Correct Answer: BDE

Explanation: B: By carrying SCSI commands over IP networks, the iSCSI protocol enables you to access block devices from across the network as if they were connected to the local system. COMSTAR provides an easier way to manage these iSCSI target devices.



D: Common Multiprotocol SCSI TARget, or COMSTAR, a software framework that enables you to convert any Oracle Solaris 11 host into a SCSI target device that can be accessed over a storage network by initiator hosts.

E: One IP port can handle multiple iSCSI target devices.

Reference: Oracle Solaris Administration: Devices and File Systems, Configuring Storage Devices With COMSTAR

QUESTION 12

Before booting testzone, 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 C
- D. zlogin -z testzone -C
- E. zlogin -C testzone F zoneadm -z testzone -C

Correct Answer: E

QUESTION 13

You need to migrate a UFS file system named /production_ufs to a ZFS file system named /production_zfs. The /production_ufs file system cannot be taken down or be out of production during the migration, and the current /production_ufs file system must remain active until the /production_zfs file system is copied and ready.

Which method allows you to meet both requirements?

1.

Copy live data from /production_ufs to /production_zfs while /production_ufs is in use.

2.

When the copy is complete, /production_zfs will contain an up-to-date copy of /production_ufs

- A. Create a snapshot of the UFS file system. Create the new ZFS file system. Use cpio to copy data from the snapshot to the new ZFS file system.
- B. Create a new Boot Environment. Create the ZFS file system. Use lucreate -m to copy data from the Current UFS file system to the new ZFS file system.
- C. Mirror the existing UFS file system by using SVM. After both submissions are in sync, migrate one of the submissions to a ZFS file system by using Live Upgrade.
- D. Create the new ZFS file system by using zfs create import to import data from the existing UFS file system into the new ZFS file system
- E. Create the new zfs file system by using the zfs create -o shadow.



Correct Answer: E

Explanation: Migrating Data With ZFS Shadow Migration ZFS shadow migration is a tool you can use to migrate data from an existing file system to a new file system. A shadow file system is created that pulls data from the original source as necessary.

You can use the shadow migration feature to migrate file systems as follows:

*

A local or remote ZFS file system to a target ZFS file system

*

A local or remote UFS file system to a target ZFS file system

Shadow migration is a process that pulls the data to be migrated:

*

Create an empty ZFS file system.

*

Set the shadow property on an empty ZFS file system, which is the target (or shadow) file system, to point to the file system to be migrated.

For example:

```
# zfs create -o shadow=nfs://system/export/home/ufsddata users/home/shadow2
```

*

Data from file system to be migrated is copied over to the shadow file system.

Reference: Transitioning From Oracle Solaris 10 to Oracle Solaris 11, Migrating Data With ZFS Shadow Migration

QUESTION 14

You are creating a non-global zone on your system.

Which option assigns a zpool to a non-global zone, and gives the zone administrator permission to create zfs file system in that zpool?

A. While creating the non-global zone, make the following entry: add device set match=/dev/rdisk/c4t0d0 end Boot the zone and log in the zone as root. Create the zpool:

```
zpool create pool2 c4t0d0
```

In the non-global zone, root can now create ZFS file system in the pool2 zpool

B. In the global zone, create the zpool: global# zpool create pool2 c4t1d0 While creating the no-global zone, make the following entry: add dataset set name=pool2 end add fs set dir=pool1 set special=pool1 set type=zfspool1 end Boot the zone, log in the zone as root, and create the zfs file system in the pool2 zpool.



C. In the global zone, create the zpool: `global#zpool create pool2 c4t1d0` While creating the global zone, make the following entry: `add dataset set name=pool2 end` Boot the zone, log in to the zone as root and create the zfs file systems in the pool2 zpool.

D. In the global zone, create the zpool and the ZFS file systems that you want to use in the non-global zone:
`global#zpool create pool2 c4t1d0` `global#zfs create pool2/data` While creating the non-global zone, make the following entry for each ZFS file system that you want to make available in the zone: `add fs set dir=/data set special=pool2/data set type=zfs end`

E. Create the zpool in the global zone: `global#zpool create pool2 c4t1d0` Boot the non-global zone, log in to the zone as root, and issue this command to delegate ZFS permissions to root: `non-global zone# zfs allow root create , destroy, mount pool2` Log in to the non-global zone create ZFS file systems in the pool2 zpool.

Correct Answer: C

Reference: Oracle Solaris ZFS Administration Guide, Adding ZFS File Systems to a Non- Global Zone

QUESTION 15

Review the ZFS dataset output that is displayed on your system:

```
M   F   /data/file5
-   F   /data/file1
R   F   /data/file3 -> /data/file13
+   F   /data/file4
```

Which four correctly describe the output?

- A. /data/file4 has been added.
- B. The link /data/file3 has been added.
- C. /data/file3 has been renamed to /data/file13.
- D. /data/file4 has been modified and is now larger.
- E. /data/file1 has been deleted.
- F. /data/file1 has been modified and is now smaller.
- G. /data/file5 has been modified.
- H. /data/file3 (a link) has been removed.

Correct Answer: ACEG

Explanation: A: + Indicates the file/directory was added in the later dataset

C: R Indicates the file/directory was renamed in the later dataset

E: - Indicates the file/directory was removed in the later dataset

G: M Indicates the file/directory was modified in the later dataset



Note: Identifying ZFS Snapshot Differences (zfs diff) You can determine ZFS snapshot differences by using the zfs diff command.

The following table summarizes the file or directory changes that are identified by the zfs diff command.

File or Directory Change Identifier

*

File or directory is modified or file or directory link changed M

*

File or directory is present in the older snapshot but not in the newer snapshot -

*

File or directory is present in the newer snapshot but not in the older snapshot +

*

File or directory is renamed R

[Latest 1Z0-820 Dumps](#)

[1Z0-820 Practice Test](#)

[1Z0-820 Exam Questions](#)