



# 1Z0-820<sup>Q&As</sup>

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

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 2

Which three options are valid methods of installing a Solaris 10 branded zone on a system running Oracle Solaris 11?

- A. Use the V2V process to migrate an existing Solaris 8 or 9 non-global zone from a Solaris 10 system to a solaris10 branded zone.
- B. Use the V2V process to migrate an existing Solaris 10 non global whole root zone from a Solaris 10 system to a solaris10 branded whole root zone.
- C. Install a solaris10 branded zone directly from the Oracle Solaris 10 media.
- D. Migrate an existing 64-bit Solaris 10 system to a solaris10 branded non-global zone using the P2V process.
- E. Migrate an existing 32 bit Solaris10 system to a solaris10 branded non-global zone using the P2V process.
- F. Use the V2V process to migrate an existing Solaris 10 non-global sparse root zone from a Solaris 10 system to a solaris10 branded sparse root zone.

Correct Answer: BDE

Explanation: B: Due to change in package system (SRV4 to IPS) there is no direct upgrade from S10 to S11 one can



use

\*

P2V converting s10 physical system to solaris10 branded zone in s11 (32-bit or 64-bit)

\*

V2V converting s10 native full root zone to solaris10 branded zone in s11 (B)

---

### QUESTION 3

You create a flash archive of the Solaris 10 global zone on the server named sysA and you name the archive s10-system.flar.

You upgrade the operating system on sysA from Solaris 10 to Oracle Solaris 11.

On sysA, you want to create a Solaris 10 non-global zone and import the flash archive into that zone. The flash archive is available on sysA via NFS through the mountpoint `/net/backup_server/backup_server`.

Which two options do you choose to create the Solaris 10 zone capable of supporting the Solaris 10 flash archive? (Choose two.)

- A. `zonecfg -z s10-zone create -t SYSSolaris10`
- B. `zonecfg -z s10-zone create setbrand=solaris10`
- C. `zonecfg -z solaris10 -t SYSSolaris10`
- D. `zonecfg -z s10-zone install -f /net/backup_server/s10-system.flar`

Correct Answer: CD

---

### QUESTION 4

You are planning the configuration for IPMP (IP Multipathing).

Best practice dictates that you must ensure that there is at least one neighbor on each subnet that responds to ICMP echo request probes to detect link failure. Which two of the following statements are true regarding the link failure detection process? (Choose two.)

- A. By default, `in.mpathd` selects the default router as the only test partner.
- B. The default failure detection time is 60 seconds (12 consecutive probes).
- C. Without test partners, `in.mpathd` cannot detect the failure of a subinterface.
- D. Arbitrary hosts on the subnet are detected by sending an "all hosts" (224.0.0.1) multicast packet.
- E. Arbitrary hosts on the subnet are detected by sending an "all routers" (224.0.0.2) multicast packet.
- F. Arbitrary hosts on the subnet are detected by examining the ARP cache (Internet-to-MAC address translation table).



Correct Answer: BD

## QUESTION 5

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

You department's backup policy is to perform a full backup to a remote system disk on Saturday.

On each weekday, 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.

Following your company policy, which two describe the correct procedure to be performed on each weekday for backing up the /data file system to a remote disk named /remote/backup?

- A. Remove the previous daily snapshot. `zfs snapshot pool1.data@daily` `zfs send -i pool1/data@sat pool1/data@daily > /remote/backup/full`
- B. Remove the previous daily snapshot. `zfs snapshot pool1.data@daily` `zfs send -i pool1/data#sat pool1/data@daily | zfs recv remote /backup/`date +%m%d%``
- C. Remove the previous daily snapshot. `zfs snapshot pool1.data@daily` `zfs send -i pool1/data@sat pool1/data@daily > /remote/backup/daily`
- D. Remove the previous daily snapshot. `zfs create -i pool1/data@sat pool1/data@daily` `zfs send pool1/data@daily | zfs remote/backup`

Correct Answer: BC

Explanation: You can use the `zfs send` command to send a copy of a snapshot stream and receive the snapshot stream in another pool on the same system or in another pool on a different system that is used to store backup data.

You can send incremental data by using the `zfs send -i` option.

For example:



host1# zfs send -i tank/dana@snap1 tank/dana@snap2 | zfs recv newtank/dana Note that the first argument (snap1) is the earlier snapshot and the second argument (snap2) is the later snapshot. In this case, the newtank/dana file system must already exist for the incremental receive to be successful. The incremental snap1 source can be specified as the last component of the snapshot name. This shortcut means you only have to specify the name after the @ sign for snap1, which is assumed to be from the same file system as snap2. For example:

host1# zfs send -i snap1 tank/dana@snap2 > zfs recv newtank/dana This shortcut syntax is equivalent to the incremental syntax in the preceding example.

Reference: Oracle Solaris ZFS Administration Guide, Sending and Receiving ZFS Data

[Latest 1Z0-820 Dumps](#)

[1Z0-820 PDF Dumps](#)

[1Z0-820 Practice Test](#)