

Question: 1

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: C, E

Question: 2

Review the following output from the `zpool status` command:

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

        NAME                        STATE        READ  WRITE CKSUM
pool1
  mirror-0      ONLINE            0      0     0
    c3t3d0      ONLINE            0      0     0
    c3t4d0      ONLINE            0      0     0
  mirror-1      ONLINE            0      0     0
    c3t5d0      ONLINE            0      0     0
    c3t6d0      ONLINE            0      0     0

errors: No known data errors
```

Which three are true for pool1?

- A. If this mirror is split, the new pool will contain disks c3t5d0 and c3t6d0.
- B. If this mirror is split, by default the new pool will contain disks c3t3d0 and c3t5d0.
- C. Data is striped across mirror-0 and mirror-1.
- D. mirror-1 is a mirrored copy of data that is stored on mirror-0.
- E. Disk c3t3d0 is a mirrored copy of disk c3c4d0.
- F. If this mirror is split, pool1 will no longer be mirrored.

Answer: B, D, E

Question: 3

The `zfs holds` command displays the following information:

NAME	TAG	TIMESTAMP
pool12/data@nov	keep	Wed May 30 12:15:12 2012

Which two statements are true?

- A. Use `zfs destroy -d pool12/data@nov` to destroy the snapshot immediately.

- C. Attempts to destroy the pool12/data@nov snapshot will not destroy the snapshot immediately.
- D. The zfs directory -R pool12/data command will destroy the file system immediately.
- E. The defer_destroy property is set to on for the pool12/data@nov data set.
- F. The userrefs property is set to 1 (or higher) for the pool12/data@nov data set.

Answer: C, E

Question: 4

Consider the following ZFS configuration:

NAME	USED	AVAIL	REFER	MOUNTPOINT
rpool/export	436M	9.21G	284M	/export
rpool/export/home	152M	9.21G	35K	/export/home
rpool/export/home/curly	35K	9.21G	35K	/export/home/curly

You have created snapshots of the home directories which are as follows:

NAME	USED	AVAIL	REFER	MOUNTPOINT
rpool/export/home@11.28.12	0	-	32K	-
rpool/export/home/curly@11.28.12	0	-	802K	-

You have another storage pool named bpool on the same system. You use the following command to store the snapshots in this pool:

```
# zfs send rpool/export/home@11.28.12 | zfs recv -f bpool@11.28.12
```

What will be created in the pool bpool as a result of this operation?

- A. bpool/export/home/
bpool/export/home/curly
- B. bpool/export/home/
bpool/export/home/curly
bpool/export/home@12.25.11
- C. bpool/export/home/
bpool/export/home/curly
bpool@12.25.11
- D. bpool/curly
bpool/curly@12.25.11
- E. bpool/curly
- F. bpool/curly
bpool@12.25.11

Answer: A

Question: 5

Consider the following commands on a newly installed system:

```
zfs set compression=on rpool
```

```
zfs get -H -O source compression rpool
```

What is the output of the second command?

- A. default
- B. -
- C. local
- D. on

Explanation:

Reference:

<http://docs.oracle.com/cd/E19082-01/817-2271/gazuk/> (querying ZFS properties for scripting)

Question: 6

You want to create a ZFS file system with the following specifications:

- lzjb compression enabled
- Cannot consume more than 2 GB from the storage pool
- Redundant data at the block level eliminated
- Mounted as /data

Which command creates the desired file system?

- A. `zfs create -o mountpoint=/data,compression=on,algorithm=lzjb,deduplication=on,quota=2g /pool1/data`
- B. `zfs create -o mountpoint=/data compression=on algorithm=lzjb deduplication=on quota=2g /pool1/data`
- C. `zfs create -o mountpoint=/data -o compression=on -o dedup=on -o quota=2g /pool1/data`
- D. `zfs create -o mountpoint=/data -o compression=on -o algorithm=lzjb -o deduplication=on -o quota=2g /pool1/data`
- E. `zfs create pool/data zfs set mountpoint=/data,quota=2g, dedup=on,compression=on /pool1/data`

Answer: D

Question: 7

Which two zpool subcommands will permanently remove a submirror from active storage pool?

- A. `remove`
- B. `detach`
- C. `destroy`
- D. `offline`
- E. `replace`
- F. `split`
- G. zpool does not permit this operation on an active storage pool unless the submirror faults.

Answer: A, B

Question: 8

Your task is to configure storage for an Oracle Solaris 11 system to support multiple web servers. Each web server will be contained in a separate zone. The system has an attached disk array configured as a JBOD (Just a Bunch Of Disks). The system also has an internal solid-state drive. The data accessed through the websites will be primarily read-only. The web servers are expected to be very busy, so configure the storage for maximum performance. Because the data is primarily static, but redundancy is required to maintain high availability in the event of a hardware failure. Data does not change often, but it is expected that the same data will be accessed many times throughout the day. Which configuration option best meets the data storage requirements?

- A. a raid2 storage pool with a separate log device

- C. a mirrored storage pool with a separate log device
- D. a three disk striped storage pool with a separate cache device
- E. a raidz1 storage pool with a separate log and cache device

Answer: B

Question: 9

The zpool configuration on serverA is:

pool 1

c3t2d0

c3t3d0

pool 2

c3t4d0

c3t5d0

The zpool configuration on serverB is:

pool1

mirror-0

c3t2d0

c3t3d0

mirror-1

c3t4d0

c3t5d0

Which option will modify the configuration on serverA to match serverB?

A. zpool destroy pool2

zpool attach pool1 c3t4d0 c3t5d0

B. zpool destroy pool2

zpool attach pool1 c3t2d0 c3t2d0 c3t4d0 c3t5d0

C. zpool destroy pool2

zpool add pool1 c3t4d0 c3t5d0

D. zpool destroy pool2

zpool mirror pool1 pool2

E. zpool destroy pool2

zpool attach pool1 c3t2d0 attach pool1 c3t3d0

zpool attach pool1 c3t4d0 attach pool1 c3t5d0

F. zpool destroy pool1

zpool destroy pool2

zpool destroy pool1 mirror c3t2d0 c3t3d0 c3t4d0 c3t5d0

Answer: F

Question: 10

Which is the result of the following command?

```
# zfs send -i dpool/sales/qrrreports@qrrreport dpool/sales/qrrreports@nth3qrrreport
```

A. An error message will be sent to standard error.

B. The dpool/sales/qrrreports@qrrreport snapshot is saved to disk.

C. The dpool/sales/qrrreports@nth3qrrreport snapshot is saved to disk

D. The difference between the First snapshot and the second snapshot will be written to disk

Question: 11

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:

NAME	USED	AVAIL	REFER	MOUNTPOINT
dpool/export	407M	130G	32K	/export
dpool/export/data	407M	130G	54.5K	/export/data
dpool/export/home/CID	226M	130G	226M	/export/home/CID
dpool/export/home/RID	180M	130G	180M	/export/home/RID
dpool/export/CID	24K	130G	226M	/export/CID

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: B**Question: 12**

To reduce the use at storage space on your server, you want 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 on write operations?

- A. zfs create -o compression=on pool1/data
- B. zpool create -o deduplication=on pool1 ; zfs create pool1/data
- C. zpool create -o dedupratio=on pool1 ; zfs create pool1/data
- D. zfs create -o dedupratio=2 pool1/data
- E. zfs create -o dedup=on pool1/data

Answer: E**Question: 13**

Which option lists default checkpoints for building an image using the Distribution Constructor?

- A. manifest-valid and ba-init
- B. ba-arch and grub-setup
- C. transfer-ips-install and pre-pkg-img-mod
- D. pkg-img mod and create-usb

Answer: A**Question: 14**

Which two statements describe the capabilities of the Distribution Constructor?

- A. ISO images for use with the Automated Installer (AI) can be created.
- B. Bootable USB images can be created for SPARC and x86 architectures.
- C. A single installation server can be used to create ISO imager, for SPARC and x86 architectures.
- D. Checkpoints are used to pause the build, thereby allowing the running of a script to modify the resulting ISO image.
- E. A single installation server can be used to create ISO images for Solaris 10 and Solaris 11.0 operating systems.

Answer: B, E
