Free Questions for 1Z0-820 by certscare

Shared by Holden on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

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?

Options:

A- While creating the non-global zone, make the following entry:

add device

set match=/dev/rdsk/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.

Question 2

The current ZFS configuration on server is:


```
pool1/data@now                      0        -      31K    -
pool1/data@monday                   0        -      31K    -
rpool/ROOT/solaris@install        280M      -    3.40G    -

pool1   3.97G    200K   3.97G    0%   1.00x   ONLINE    -
pool3   7.94G   6.25G   1.69G   78%   1.00x   ONLINE    -
rpool   15.9G   11.6G   4.24G   73%   1.00x   ONLINE    -

pool1                             200K   3.91G    31K   /pool1
pool1/data                         31K   3.91G    31K   /data
pool1/data2                        18K   3.91G    31K   /data2
pool3                             6.25G   1.56G    32K   /pool3
pool3/IPS                         6.25G   1.56G   6.25G   /pool3/IPS
```

You need to backup the /data file system while the file system is active.

Select the option that creates a full backup of the /data file system and stores the backup on server in the pool named backup.

Options:

A- Mount -F nfs system: /backup / mnt

zfs snapshot pool/data@monday>/mnt/Monday

B- Mount -F nfs systemB: /backup/mnt

zfs snapshot pool1/data@Monday

zfs clone pool1/data@monday/mnt/Monday

C- Zfs send pool1/data@Monday | ssh system zfs recv backup/monday

D- Zfs snapshot pool1/data@Monday | ssh system zfs recv backup/monday

Answer:

D

Explanation:

Example (assuming there exists a snapshot 'tank/test@1).

Full backup

Now let's do a full initial backup from the 'tank/test@1 snapshot:

```
# zfs send tank/test@1 | zfs receive tank/testback
```

Question 3

Question Type: MultipleChoice

You created a new zpool. Now you need to migrate the existing ZFS file system from pool1/prod to pool2/prod.

You have these requirements:

1. Users must have access to the data during the migration, so you cannot shutdown the file system while the migration takes place.

2. Because you want to copy the data as quickly as possible, you need to increase the server resources devoted to the ZFS migration.

Which method would you use to modify the ZFS shadow migration daemon defaults to increase the concurrency and overall speed of migration?

Options:

A- Svccfg - s filesystem/shadowd:default
setprop config_params/shadow_threads=integer: 16
end

svcadm refresh filesystem/shadowd: default

B- Specify the -b <blocksize> option with the zfs create command and increase the value of <blocksize>

C- Use the -o -volblocksize=<blocksize>option with the zfs create command and increase the value of the default <blocksize>.

D- Svccfg -s filesystem/zfs: default
setprop config_params/shadow_threads = integer: 16
end
svcadm refresh filesystem/zfs:default

Answer:

A

Explanation:

shadowd is a daemon that provides background worker threads to migrate data for a shadow migration. A shadow migration gradually moves data from a source file system into a new "shadow" file system. Users can access and change their data within the shadow file system while migration is occurring.

The shadowd service is managed by the service management facility, smf(5). Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). The service's status can be queried using the svcs(1) command.

The svccfg(1M) command can be used to manage the following parameter related to shadowd:

config_params/shadow_threads

Note: Oracle Solaris 11: In this release, you can migrate data from an old file system to a new file system while simultaneously allowing access and modification of the new file system during the migration process.

Setting the shadow property on a new ZFS file system triggers the migration of the older data. The shadow property can be set to migrate data from the local system or a remote system with either of the following values:

file:///path

nfs://host:path

Question 4

You are troubleshooting the Oracle Solaris11 Automated Installer (AI), which is not connecting with the IPS software repository.

Which three steps will help determine the cause of DNS name resolution failure?

Options:

A- Verify the contents of /etc/resolve.conf.

B- Run netstat -nr to verify the routing to the DNS server.

C- Ping the IP address of the IPS server to verify connectivity.

D- On the installation server, verify that the menu.1st file for the client points to a valid boot arc hive.

E- Run df -k to verify that the boot directory containing the boot archive is loopback mounted under /etc/netboot.

F- Run the command /sbin/dhcpinfo DNSserv to ensure that the DHCP server providing the DNS server information.

Answer:

A, B, F

Explanation:

Check DNS

* (A) Check whether DNS is configured on your client by verifying that a non-empty /etc/resolv.conf file exists.

* (F) If /etc/resolv.conf does not exist or is empty, check that your DHCP server is providing DNS server information to the client:

/sbin/dhcpinfo DNSserv

If this command returns nothing, the DHCP server is not set up to provide DNS server information to the client. Contact your DHCP administrator to correct this problem.

* (B) If an /etc/resolv.conf file exists and is properly configured, check for the following possible problems and contact your system administrator for resolution:

** The DNS server might not be resolving your IPS repository server name.

** No default route to reach the DNS server exists.

Question 5

New features wore added to ZFS in Oracle Solaris11. Your justification to upgrade from Solaris10 to oracle Solaris11 is that it will be possible to take advantage of the enhancements that were made to ZFS.

Identify the three ZFS functions and features that are included in Oracle Solaris 11, but not in Solaris 10.

Options:

A- Encrypted ZFS datasets

B- Ability for ZFS to detect and remove redundant data from the tile system

C- Shadow Data Migration

D- Ability to split a mirrored ZFS storage pool

E- Ability to use ZFS on the boot drive and boot to a ZFS root file system.

F- elimination of the swap file system when using ZFS on the root disk

Answer:

A, B, E

Explanation:

A: ZFS encryption was introduced in ZFS Pool Version Number 30, Solaris Nevada b149.

Filesystem encryption since Solaris 11 Express

B: Deduplication was introduced in ZFS Pool Version Number 21, Solaris Nevada b128.

E: Booting From a ZFS Root File System

Both SPARC based and x86 based systems use the new style of booting with a boot archive, which is a file system image that contains the files required for booting. When a system is booted from a ZFS root file system, the path names of both the boot archive and the kernel file are resolved in the root file system that is selected for booting.

Question 6

You are setting up a local IPS package repository on your Oracle Solaris11 server: solaris.example.com.

You want to point the existing local IPS publisher to the new local IPS repository located in /repo.

These are the stops that you have followed:

1. Download and rsync the contents of the Oracle Solaris11 repository ISO image to the /repo directory.

2. Configure the repository server service properties. The svcprop command display, the IPS related properties:

pkg/inst_root astring/repo

pkg/readonly Boolean true

The 1s command displays the contents of the /repo directory:

#ls/repo

Pkg5.repository publisher

The svcs publisher command shows the svc: /application/pkg/server: default service is online.

The pkg publisher command shows the svc: /application/pkg/server: default service is online.

The pkg publisher command still displays:

PUBLISHER TYPE STATUS URI

solaris origin online http://pkg.oracle.com/solaris/release/

Which steps needs to be performed to set the local IPS publisher to the local IPS repository/repo?

Options:

A- Issue the pkgrepo refresh -s command to refresh the repository.

B- Restart the svc:/application/pkg/server:default service.

C- pkg set-publisher command to set the new repository location.

D- Issue the pkgrepo rebuild command to rebuild the repository.

E- Issue the pkgrepo set command to set the new repository location.

Answer:

C

Explanation:

Set the Publisher Origin To the File Repository URI

To enable client systems to get packages from your local file repository, you need to reset the origin for the solaris publisher. Execute the following command on each client:

Example:

pkg set-publisher -G '*' -M '*' -g /net/host1/export/repoSolaris11/ solaris

Question 7

United States of America export laws include restrictions on cryptography.

Identify the two methods with which these restrictions are accommodated in the Oracle Solaris 11 Cryptographic Framework.

Options:

A- Corporations must utilize signed X.509 v3 certificates.

B- A third-party provider object must be signed with a certificate issued by Oracle.

C- Loadable kernel software modules must register using the Cryptographic Framework SPI.

D- Third-party providers must utilize X.509 v3 certificates signed by trusted Root Certification Authorities.

E- Systems destined for embargoed countries utilize loadable kernel software modules that restrict encryption to 64 bit keys.

Answer:

B, C

Explanation:

B: Binary Signatures for Third-Party Software

The elfsign command provides a means to sign providers to be used with the Oracle Solaris Cryptographic Framework. Typically, this command is run by the developer of a provider.

The elfsign command has subcommands to request a certificate from Sun and to sign binaries. Another subcommand verifies the signature. Unsigned binaries cannot be used by the Oracle Solaris Cryptographic Framework. To sign one or more providers requires the certificate from Sun and the private key that was used to request the certificate.

C: Export law in the United States requires that the use of open cryptographic interfaces be restricted. The Oracle Solaris Cryptographic Framework satisfies the current law by requiring that kernel cryptographic providers and PKCS #11 cryptographic providers be signed.

Question 8

Question Type: MultipleChoice

Which modification needs to be made to the Service Management Facility before you publish a new package to the IPS repository?

Options:

A- The pkg.depotd must be disabled.

B- The pkg/readonly property for the application/pkg/server service must be set to false

C- The pkg/writable_root property for the application/pkg/server service must be set to true.

D- The pkg/image_root property for the application/pkg/server service must be set to the location of the repository.

Answer:

D

Explanation:

pkg/image_root

(astring) The path to the image whose file information will be used as a cache for file data.

Incorrect answers:

A: pkg.depotdis the depot server for the image packaging system. It provides network access to the data contained within a package repository. Clients that do not support direct access to a repository through the file system, or for which network access is the only available or preferred method of transport, typically use the package depot.

B: pkg/readonly(boolean)

Sets whether modifying operations, such as those initiated by pkgsend(1), are disabled. Retrieval operations are still available. This property cannot be true when the pkg/mirror property is true. The default value is true.

C: pkg/writable_root

(astring) The file system path to a directory to which the program has write access. This is used with the-readonlyoption to enable the depot server to create files, such as search indexes, without needing write access to the package information.

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
 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 c4t0d0and c4t2d0 accessible under the dev_data mount point.

Options:

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- zpool 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

Answer:

D

Explanation:

In this Solaris release, you can use thezpool splitcommand to split a mirrored storage pool, which detaches a disk or disks in the original mirrored pool to create another identical pool.

After the split operation, import the new pool.