Free Questions for 1Z0-820

Shared by Knight on 04-10-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

View the Exhibit to inspect the file system configuration on your server.

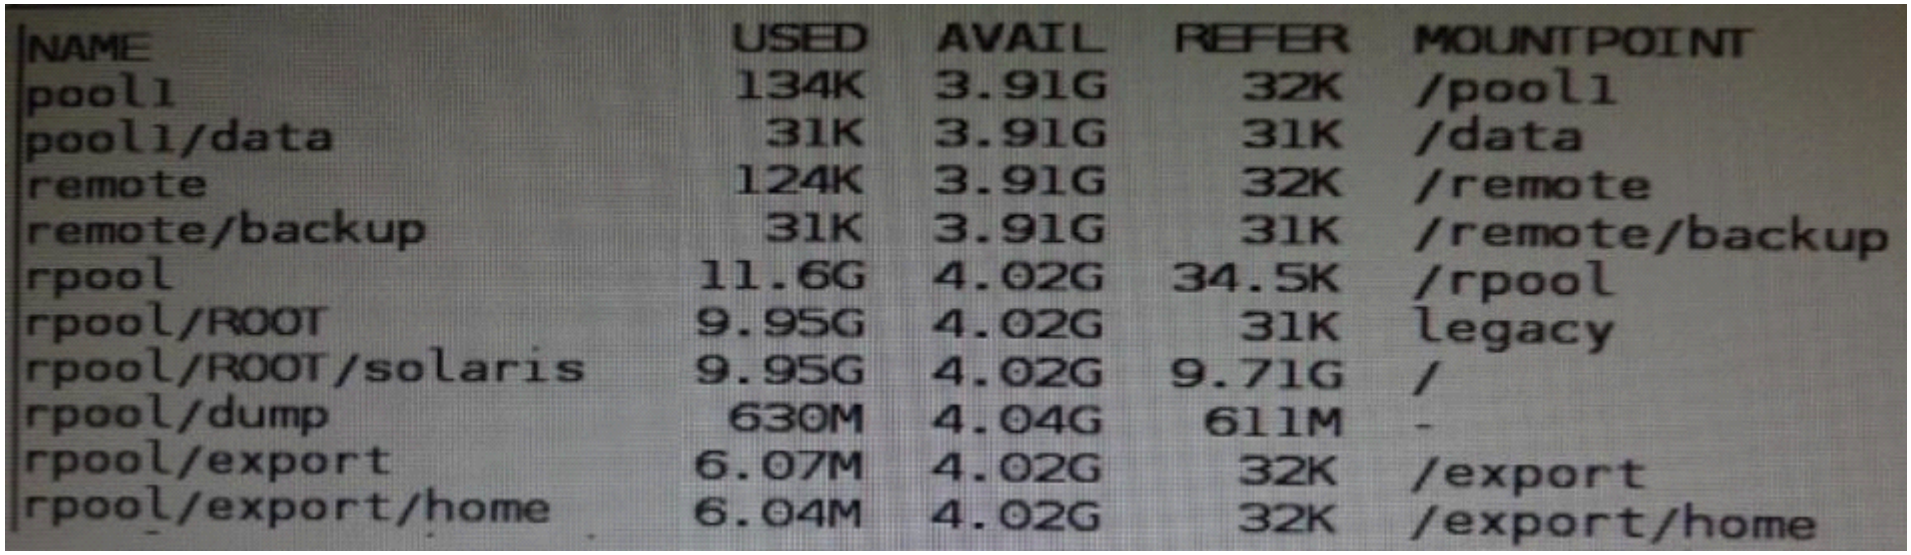

NAME	USED	AVAIL	REFER	MOUNTPOINT
pool1	134K	3.91G	32K	/pool1
pool1/data	31K	3.91G	31K	/data
remote	124K	3.91G	32K	/remote
remote/backup	31K	3.91G	31K	/remote/backup
rpool	11.6G	4.02G	34.5K	/rpool
rpool/ROOT	9.95G	4.02G	31K	legacy
rpool/ROOT/solaris	9.95G	4.02G	9.71G	/
rpool/dump	630M	4.04G	611M	-
rpool/export	6.07M	4.02G	32K	/export
rpool/export/home	6.04M	4.02G	32K	/export/home

Your department's backup policy is to perform a full backup to 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 Saturday for backing tip the /data file system to a remote disk named /remote/backup?

Options:

A- On Saturday:
zfs snapshot pool1/data@sat
zfs send pool1/data@sat > /remote/backup/full

B- On Saturday:
zfs create snapshotpool1/data@sat
zfs send pool1/data@sat | zfs recv remote/backup/ 'date' '+%m%d%y'

C- On Saturday:
zfs create snapshotpool1/data@sat
zfs send pool1/data@sat > /remote/backup/full

D- On Saturday:
zfs create snapshotpool1/data@sat
zfs send pool1/data@sat | zfsrecv remote/backup

Answer:

B, C

Explanation:

You can use thezfs sendcommand 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 thezfs send-ioption.

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.

Question 2

Question Type: MultipleChoice

You need to migrate a UFS file system named /production_ufs to a ZFS file system named /production_ufs. 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 /ptoduction_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

Options:

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.

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/ufsdata users/home/shadow2

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

Incorrect answers:

B: lucreate-- create a new boot environment

Note: ZFS is the default root file system.

UFS is a supported legacy file system, but it is not supported as a bootable root file system.

Question 3

Question Type: MultipleChoice

After installing the OS, you boot the system and notice that the syslogd daemon is not accepting messages from remote systems.

Which two options should you select to modify the syslogd daemon configuration so that it accepts messages from remote systems?

Options:

A- svccfg -s svc:/system/system -log setprop start/exec= "syslogd -t"
Restart the syslogd daemon.

B- Set the following parameter in the /etc/syslogd.conf file: LOG_FROM_REMOTE= YES
Restart the syslogd daemon.

C- svcadm enable svc:/system/system -log/config/log_from_remote
Restart the syslogd daemon.

D- svccfg -s svc:/system/system-log setprop config/log_from_remote=true
Restart the syslogd daemon.

E- Set the following parameter in the /etc/default/syslogd file: LOG_FROM_REMOTE=YES
Restart the syslogd daemon.

Answer:

B, D

Explanation:

B: The /etc/default/syslogd file contains the following default parameter settings. See FILES.

LOG_FROM_REMOTE

Specifies whether remote messages are logged. LOG_FROM_REMOTE=NO is equivalent to the -t command-line option. The default value for LOG_FROM_REMOTE is YES.

Question 4

Question Type: MultipleChoice

The crash dump notification on your server is:

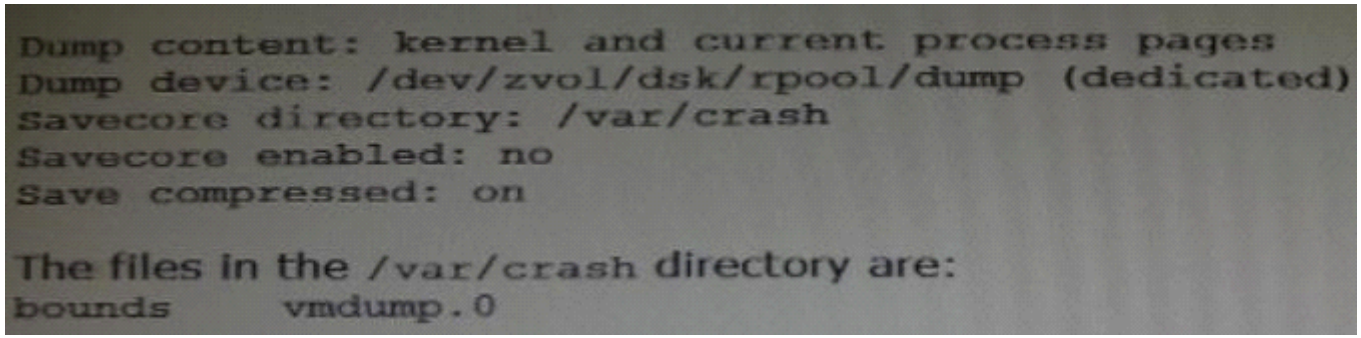

Documentation states that there should be two core files for each crash dump in the /var/crash directory named vmdump.0

Which command should you choose to display theses two files?

Options:

A- savecore -f vmdump.0

B- dumpadm uncompressed

C- gunzip vmdump.0

D- dumpadm -z off

Answer:

A

Explanation:

Decompress using savecore -f vmdump.0

savecore - save a crash dump of the operating system

-f dumpfile Attempt to save a crash dump from the speci-

fied file instead of from the system's

current dump device. This option may be use-

ful if the information stored on the dump

device has been copied to an on-disk file by

means of the dd(1M) command.

Question 5

Question Type: MultipleChoice

View the exhibit to inspect the file system configuration on your server.

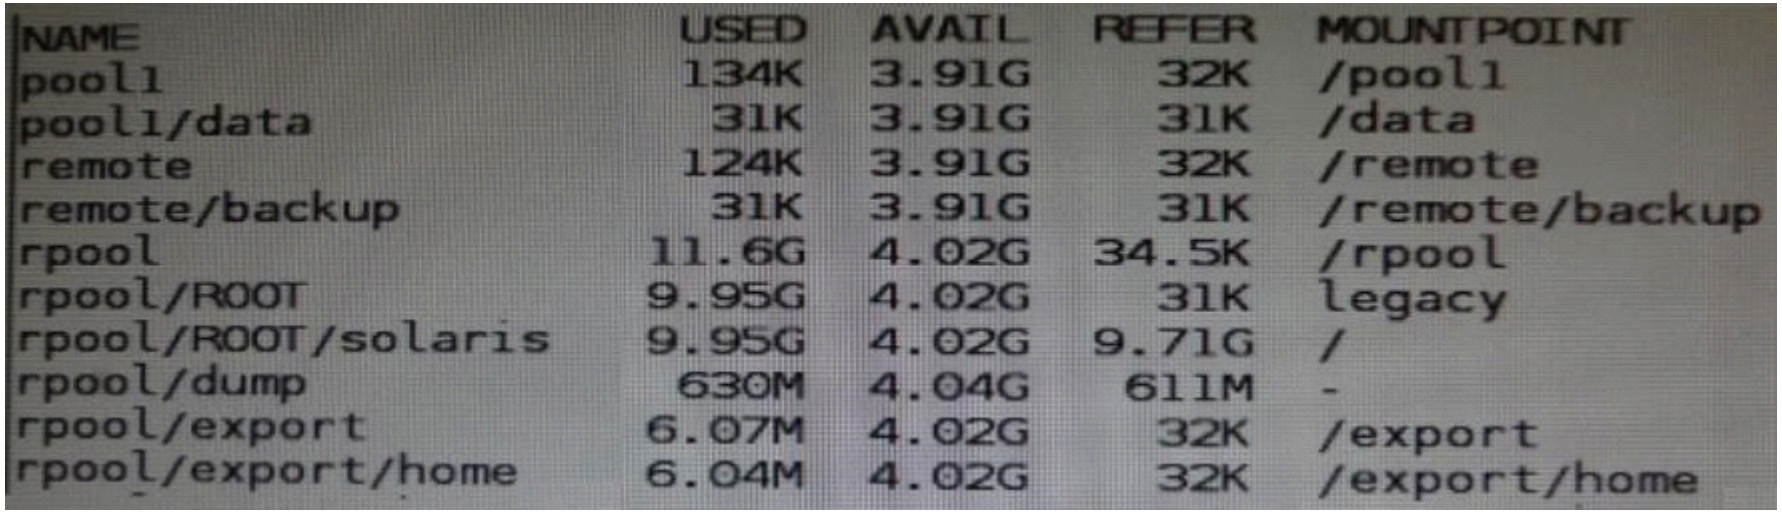

NAME	USED	AVAIL	REFER	MOUNTPOINT
pool1	134K	3.91G	32K	/pool1
pool1/data	31K	3.91G	31K	/data
remote	124K	3.91G	32K	/remote
remote/backup	31K	3.91G	31K	/remote/backup
rpool	11.6G	4.02G	34.5K	/rpool
rpool/ROOT	9.95G	4.02G	31K	legacy
rpool/ROOT/solaris	9.95G	4.02G	9.71G	/
rpool/dump	630M	4.04G	611M	-
rpool/export	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?

Options:

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

Answer:

B, C

Explanation:

You can use thezfs sendcommand 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 thezfs send-ioption.

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.

Question 6

Question Type: MultipleChoice

You are setting up a local image packaging System (IPS) package repository on your Oracle Solaris 11 server. The information and specifications that you have are as follows:

The Oracle Solaris11 repository ISO image has been downloaded into the /repo directory (a zfs file system).

The current publisher is:

PUBLISHER TYPE STATUS URI

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

You will be replacing the current publisher with:

PUBLISHER TYPE STATUS URI

solaris origin online http://solaris.example.com/

The location of the repository will be /export/IPS. This ZFS file system has already been created.

Among the steps you will perform is to set the publisher to the local repository by using the pkg set - publisher command.

Which six other steps are required to set up the local IPS package repository?

Options:

A- Perform an rsync on the ISO image to copy the files from the ISO image to the /export/IPS file system.

B- chmod 700 on/export/IPS

C- Use the svccfg command to set the pkg/inst_root property to export/IPS.

D- Use the svccfg command to set the pkg.inst_root property to /export/IPS.

E- Use the svccfg command to set the pkg/readonly property to the application/pkg/server service to true.

F- Use the svccfg command to set the pkg/readonly property for the application/pkg/server service false.

G- Refresh the application/pkg/server service with the svcadm refresh command.

H- Refresh the package repository with the pkgrep refresh command.

I- Enable the application /pkg/server service.

J- Run the pkhrepo rebuild command to rebuild the repository catalog.

Answer:

A, B, D, E, G, I

Explanation:

B: Set the correct permissions.

Serving a Local Repository Using SMF

First, create a ZFS file system to hold the repository.

Next, make a local copy of the IPS repository file.

Make the contents of the repository file available to thepkg.depotd(1M) server, using thelofiadmcommand to mount the ISO image.

4. (A) Copy the repository files to the ZFS file system you created. This will increase the performance of repository accesses and avoid the need to remount the .iso image each time the system restarts. (Be sure to use /mnt/repo, not /mnt/repo/, so that you copy the repo directory and not just the files and subdirectories in the repo directory.)

rsync -aP /mnt/repo /export/repo2010_11

You can use the df command to confirm the copy:

5. Once the files are copied, unmount the image and deallocate the block device.

6. Now that your local copy is ready, use thesvccfgcommand to configure the repository server service, specifying the location of your local repository and settingreadonlytotrue:

(D)# svccfg -s application/pkg/server setprop pkg/inst_root=/export/repo2010_11/repo

(E)# svccfg -s application/pkg/server setprop pkg/readonly=true

7. Start thepkg.depotdrepository service:

(G) # svcadm refresh application/pkg/server

(I) # svcadm enable application/pkg/server

http://www.oracle.com/technetwork/articles/servers-storage-admin/localrepositories-1377242.html

Question 7

Question Type: MultipleChoice

View the Exhibit.

```
<?xml version="1.0" encoding="UTF-8"?>
<!--

 Copyright (c) 2011, Oracle and/or its affiliates. All rights reserved.

-->
<!DOCTYPE auto_install SYSTEM "file:///usr/share/install/ai.dtd.1">

<auto_install>
    <ai_instance name="zone_default">
        <target>
            <logical>
                <zpool name="rpool">
                    <filesystem name="export" mountpoint="/export"/>
                    <filesystem name="export/home"/>
                    <be name="solaris">
                        <options>
                            <option name="compression" value="on"/>
                        </options>
                    </be>
                </zpool>
            </logical>
        </target>

        <software type="IPS">
            <destination>
                <image>
                    <!-- Specify locales to install -->
                    <facet set="false">facet.locale.*</facet>
                    <facet set="true">facet.locale.de</facet>
                    <facet set="true">facet.locale.de_DE</facet>
                    <facet set="true">facet.locale.en</facet>
                    <facet set="true">facet.locale.en_US</facet>
                    <facet set="true">facet.locale.es</facet>
                    <facet set="true">facet.locale.es_ES</facet>
                    <facet set="true">facet.locale.fr</facet>
                    <facet set="true">facet.locale.fr_FR</facet>
                    <facet set="true">facet.locale.it</facet>
                    <facet set="true">facet.locale.it_IT</facet>
                    <facet set="true">facet.locale.ja</facet>
                    <facet set="true">facet.locale.ja_*</facet>
```

The file came from your Automated Installer (AI) install server. The file is _____.

Options:

A- an AI SC profile for non-global zones

B- the default AI config file for non-global zones

C- the default AI manifest for non-global zones

D- a custom AI manifest

Answer:

D

Explanation:

ai_manifest

- Automated installation manifest file format

Synopsis

/usr/share/install/ai.dtd.1

Some customizations have been made, such as the selection of specific locales.

Question 8

Question Type: MultipleChoice

The following information is displayed about the compress/zjp software package, which Is currently installed on this system:

NAME (PUBLISHER) VERSION IFO

Compress/zip 3.1.2-0.175.0.0.0.0.537 if-

NAME VERSION DATE COMMENT

Compress/zip 3.1 09 Dec 2011 04:50:38 EST None

Which statement describes the information that is displayed tor the compress/zip software package?

Options:

A- This package cannot be removed.

B- This package can be updated to a new version when the new version of the package becomes available.

C- This package cannot be updated.

D- This package can be updated to version 3.1.3 but not 3.2.

E- This package cannot be downgraded to version 3.1.1.

Answer:

D

Explanation:

An "f" in the F column indicates the package is frozen. If a package is frozen, you can only install or update to packages that match the frozen version.

Note: The "i" in the I column indicates that these packages are installed in this image.

References: Adding and Updating Oracle Solaris 11 Software Packages, Showing Package Install State Information

Question 9

Question Type: MultipleChoice

When you issue the "gzip: zommand 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?

Options:

A- pkg info|grep gzip

B- pkg list SUNWgzip

C- pkg contents gzip

D- pkg search gzip

Answer:

D

Explanation:

Searching for Packages

Use the pkg search command to search for packages whose data matches the specified pattern.

Like thepkg contentscommand, thepkg searchcommand examines the contents of packages. While thepkg contentscommand returns the contents, thepkg searchcommand returns the names of packages that match the query.

pkg search

search [-HIaflpr] [-o attribute ...] [-s repo_uri] query

Search for matches to the query, and display the results.

Which tokens are indexed are action-dependent, but may

include content hashes and pathnames.

Note: pkg is the retrieval client for the image packaging system. With

a valid configuration, pkg can be invoked to create locations for

packages to be installed, called 'images', and install packages

into those images. Packages are published by publishers, who may

make their packages available at one or more repositories. pkg,

then, retrieves packages from a publisher's repository and

installs them into an image.

Question 10

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
```