



Free Questions for EX200

Shared by Oliver on 20-10-2022

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)



# Question 1

---

Question Type: MultipleChoice

---

Part 2 (on Node2 Server)

Task 6 [Implementing Advanced Storage Features]

Add a new disk to your virtual machine with a size of 10 GiB

On this disk, create a VDO volume with a size of 50 GiB and mount it persistently on /vbread with xfs filesystem



Options:

---

A- Explanation:

\*

```
[root@node2 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vdd 252:48 0 5G 0 disk
vde 252:64 0 10G 0 disk
[root@node2 ~]# yum install kmod-kvdo vdo
[root@node2 ~]# systemctl enable --now vdo
[root@node2 ~]# systemctl start vdo
[root@node2 ~]# systemctl status vdo
[root@node2 ~]# vdo create --name=vdo1 --device=/dev/vde --vdoLogicalSize=50G
[root@node2 ~]# vdostats --hu
Device Size Used Available Use% Space saving%
/dev/mapper/vdo1 10.0G 4.0G 6.0G 40% N/A
[root@node2 ~]# mkfs.xfs -K /dev/mapper/vdo1
*
[root@node2 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vde 252:64 0 10G 0 disk
vdo1 253:4 0 50G 0 vdo
[root@node2 ~]# mkdir /vbread
[root@node2 ~]# blkid
/dev/mapper/vdo1: UUID='1ec7a341-6051-4aed-8a2c-4d2d61833227' BLOCK_SIZE='4096'
TYPE='xfs'
[root@node2 ~]# vim /etc/fstab
UUID=1ec7a341-6051-4aed-8a2c-4d2d61833227 /vbread xfs defaults,x-
systemd.requires=vdo.service 0 0
[root@node2 ~]# mount /dev/mapper/vdo1 /vbread/
[root@node2 ~]# df -hT
```

```
Filesystem Type Size Used Avail Use% Mounted on
/dev/mapper/vdo1 xfs 50G 390M 50G 1% /vbreadd
```

Answer:

A

## Question 2

Question Type: MultipleChoice

Part 2 (on Node2 Server)

Task 4 [Managing Logical Volumes]

Resize the logical volume, lvrz and reduce filesystem to 4600 MiB. Make sure the the filesystem contents remain intact with mount point /datarz

(Note: partitions are seldom exactly the size requested, so anything within the range of 4200MiB to 4900MiB is acceptable)

Options:

A- Explanation:

\*

```
[root@node2 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vdb 252:16 0 5G 0 disk
vdb1 252:17 0 4.2G 0 part
vgrz-lvrz 253:2 0 4.1G 0 lvm /datarz
vdc 252:32 0 5G 0 disk
vdc1 252:33 0 4.4G 0 part
datavg-datalv 253:3 0 3.9G 0 lvm /data
vdd 252:48 0 5G 0 disk
vde 252:64 0 10G 0 disk
[root@node2 ~]# lvs
LV VG Attr LSize Pool Origin Data% Meta% Move Log Cpy%Sync Convert
lvrz vgrz -wi-ao---- 4.10g
[root@node2 ~]# vgs
VG #PV #LV #SN Attr VSize VFree
vgrz 1 1 0 wz--n- <4.15g 48.00m
[root@node2 ~]# parted /dev/vdb print
Number Start End Size Type File system Flags
1 1049kB 4456MB 4455MB primary lvm
```

```
*
[root@node2 ~]# df -hT
Filesystem Type Size Used Avail Use% Mounted on
/dev/mapper/vgrz-lvrz ext4 4.0G 17M 3.8G 1% /datarz
[root@node2 ~]# parted /dev/vdb mkpart primary 4456MiB 5100MiB
[root@node2 ~]# parted /dev/vdb set 2 lvm on
[root@node2 ~]# udevadm settle
[root@node2 ~]# pvcreate /dev/vdb2
Physical volume '/dev/vdb2' successfully created.
*
[root@node2 ~]# vgextend vgrz /dev/vdb2
Volume group 'vgrz' successfully extended
[root@node2 ~]# lvextend -r -L 4600M /dev/vgrz/lvrz
Size of logical volume vgrz/lvrz changed from 4.10 GiB (1050 extents) to 4.49 GiB (1150 extents).
Logical volume vgrz/lvrz successfully resized.
[root@node2 ~]# resize2fs /dev/vgrz/lvrz
[root@node2 ~]# df -hT
Filesystem Type Size Used Avail Use% Mounted on
/dev/mapper/vgrz-lvrz ext4 4.4G 17M 4.2G 1% /datarz
```

Answer:

---

A

## Question 3

---

Question Type: MultipleChoice

---

Configure autofs to automount the home directories of LDAP users as follows:

host.domain11.example.com NFS-exports /home to your system.

This filesystem contains a pre-configured home directory for the user ldapuser11 ldapuser11's home directory is host.domain11.example.com /rhome/ldapuser11 ldapuser11's home directory should be automounted locally beneath /rhome as /rhome/ldapuser11

Home directories must be writable by their users

ldapuser11's password is 'password'.

Options:

---

A- Explanation:

vim /etc/auto.master /rhome /etc/auto.misc

```
wq!  
# vim /etc/auto.misc  
ldapuser11 --rw,sync host.domain11.example.com:/rhome/ldapuser11 :wq!  
#service autofs restart  
service autofs reload  
chkconfig autofs on  
su -ldapuser11  
Login ldapuser with home directory  
# exit
```

Answer:

A



## Question 4

Question Type: MultipleChoice

Part 1 (on Node1 Server)

Task 15 [Running Containers]

Create a container named logserver with the image rhel8/rsyslog found from the registry registry.domain15.example.com:5000

The container should run as the root less user shangril

a. use redhat as password [sudo user]

Configure the container with systemd services as the shangrila user using the service name, "container-logserver" so that it can be persistent across reboot.

Use admin as the username and admin123 as the credentials for the image registry.

Options:

A- Explanation:

\*

```
[root@workstation ~]# ssh shangrila@node1  
[shangrila@node1 ~]$ podman login registry.domain15.example.com:5000  
Username: admin  
Password:  
Login Succeeded!  
[shangrila@node1 ~]$ podman pull registry.domain15.example.com:5000/rhel8/rsyslog  
[shangrila@node1 ~]$ podman run -d --name logserver
```

```
registry.domain15.example.com:5000/rhel8/rsyslog
021b26669f39cc42b8e94eab886ba8293d6247bf68e4b0d76db2874aef284d6d
[shangrila@node1 ~]$ mkdir -p ~/.config/systemd/user
[shangrila@node1 ~]$ cd ~/.config/systemd/user
*
[shangrila@node1 user]$ podman generate systemd --name logserver --files --new
/home/shangrila/.config/systemd/user/container-logserver.service
[shangrila@node1 ~]$ systemctl --user daemon-reload
[shangrila@node1 user]$ systemctl --user enable --now container-logserver.service
[shangrila@node1 ~]$ podman ps
CONTAINER ID IMAGE COMMAND CREATED STATUS PORTS NAMES
7d9f7a8a4d63 registry.domain15.example.com:5000/rhel8/rsyslog:latest /bin/rsyslog.sh 2
seconds ago logserver
[shangrila@node1 ~]$ sudo reboot
[shangrila@node1 ~]$ cd .config/systemd/user
[shangrila@node1 user]$ systemctl --user status
```

Answer:

---

A

## Question 5

---

Question Type: MultipleChoice

---

Part 2 (on Node2 Server)

Task 5 [Managing Logical Volumes]

Add an additional swap partition of 656 MiB to your system. The swap partition should automatically mount when your system boots

Do not remove or otherwise alter any existing swap partition on your system

Options:

---

A- Explanation:

\*

```
[root@node2 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vdc 252:32 0 5G 0 disk
vdc1 252:33 0 4.1G 0 part
datavg-datalv 253:3 0 3.9G 0 lvm /data
vdd 252:48 0 5G 0 disk
```

```
vde 252:64 0 10G 0 disk
[root@node2 ~]# swapon -s
Filename Type Size Used Priority
/dev/dm-1 partition 2097148 1548 -2
[root@node2 ~]# free -m
total used free shared buff/cache available
Mem: 1816 1078 104 13 633 573
Swap: 2047 1 2046
[root@node2 ~]# parted /dev/vdc print
Number Start End Size Type File system Flags
1 1049kB 4404MB 4403MB primary lvm
*
[root@node2 ~]# parted /dev/vdc mkpart primary linux-swap 4404MiB 5060MiB
[root@node2 ~]# mkswap /dev/vdc2
Setting up swspace version 1, size = 656 MiB (687861760 bytes)
no label, UUID=9faf818f-f070-4416-82b2-21a41988a9a7
[root@node2 ~]# swapon -s
Filename Type Size Used Priority
/dev/dm-1 partition 2097148 1804 -2
[root@node2 ~]# swapon /dev/vdc2
*
[root@node2 ~]# swapon -s
Filename Type Size Used Priority
/dev/dm-1 partition 2097148 1804 -2
/dev/vdc2 partition 671740 0 -3
[root@node2 ~]# blkid
/dev/vdc2: UUID='9faf818f-f070-4416-82b2-21a41988a9a7' TYPE='swap'
PARTUUID='0f22a35f-02'
[root@node2 ~]# vim /etc/fstab
UUID=9faf818f-f070-4416-82b2-21a41988a9a7 swap swap defaults 0 0
[root@node2 ~]# reboot
[root@node2 ~]# swapon -s
Filename Type Size Used Priority
/dev/dm-1 partition 2097148 1804 -2
/dev/vdc2 partition 671740 0 -3
```

Answer:

A

## Question 6

Question Type: MultipleChoice

Part 1 (on Node1 Server)

Task 3 [Managing Local Users and Groups]

Create the following users, groups and group memberships:

A group named sharegrp

A user harry who belongs to sharegrp as a secondary group

A user natasha who also belongs to sharegrp as a secondary group

A user copper who does not have access to an interactive shell on the system and who is not a member of sharegrp.

harry, natasha and copper should have the password redhat

### Options:

---

A- Explanation:

```
* [root@node1 ~]# groupadd sharegrp
[root@node1 ~]# useradd harry
[root@node1 ~]# useradd natasha
[root@node1 ~]# usermod -aG sharegrp harry
[root@node1 ~]# usermod -aG sharegrp natasha
[root@node1 ~]# useradd -s /sbin/nologin copper
[root@node1 ~]# echo 'redhat' | passwd --stdin harry
[root@node1 ~]# echo 'redhat' | passwd --stdin natasha
[root@node1 ~]# echo 'redhat' | passwd --stdin copper
### For Checking ###
[root@node1 ~]# su - copper
This account is currently not available.
[root@node1 ~]# su - natasha
[root@node1 ~]# id
[root@node1 ~]# su - harry
[root@node1 ~]# id
```

### Answer:

---

A

## Question 7

---

Question Type: MultipleChoice

---



Part 1 (on Node1 Server)

Task 13 [Archiving and Transferring Files & SELinux]

Create a backup file named /root/backup.tar.bz2. The backup file should contain the content of /usr/local and should be zipped with bzip2 compression format.

Furthermore, ensure SELinux is in enforcing mode. If it is not, change SELinux to enforcing mode.

### Options:

---

A- Explanation:

\*

```
[root@node1 ~]# tar cvf /root/backup.tar /usr/local/
tar: Removing leading `/' from member names
/usr/local/
/usr/local/bin/
/usr/local/etc/
[root@node1 ~]# ls
backup.tar
[root@node1 ~]# file backup.tar
backup.tar: POSIX tar archive (GNU)
[root@node1 ~]# bzip2 backup.tar
[root@node1 ~]# ls
backup.tar.bz2
[root@node1 ~]# file backup.tar.bz2
backup.tar.bz2: bzip2 compressed data, block size = 900k
*
```

```
[root@node1 ~]# sestatus
SELinux status: enabled
[root@node1 ~]# cat /etc/selinux/config
SELINUX=enforcing
SELINUXTYPE=targeted
[root@node1 ~]# reboot
### For Checking ###
[root@node1 ~]# sestatus
SELinux status: enabled
```

### Answer:

---

A

## Question 8

---

Question Type: MultipleChoice

---

Configure the permissions of /var/tmp/fstab

Copy the file /etc/fstab to /var/tmp/fstab. Configure the permissions of /var/tmp/fstab so that:

the file /var/tmp/fstab is owned by the root user.

the file /var/tmp/fstab belongs to the group root.

the file /var/tmp/fstab should not be executable by anyone.

the user natasha is able to read and write /var/tmp/fstab.

the user harry can neither write nor read /var/tmp/fstab.

all other users (current or future) have the ability to read /var/tmp/fstab.

### Options:

---

A- Explanation:

```
cp -a /etc/fstab /var/tmp
```

```
cd /var/tmp
```

```
ls -l
```

```
getfacl /var/tmp/fstab
```

```
chmod ugo-x /var/tmp/fstab
```

[ No need to do this, there won't be execute permission for the file by default]

```
# setfacl -m u:natasha:rw /var/tmp/fstab # setfacl -m u:harry:0 /var/tmp/fstab(zero)
```

[Read permission will be there for all the users, by default. Check it using `ls -l /var/tmp/fstab`]

Verify by

```
[ ls -la /var/tmp/fstab]
```

### Answer:

---

A

## Question 9

---

Question Type: MultipleChoice

---

In the system, mounted the iso image /root/examine.iso to/mnt/iso directory. And enable automatically mount (permanent mount) after restart system.

### Options:

---

A- Explanation:

```
mkdir -p /mnt/iso
```

```
/etc/fstab:
```

```
/root/examine.iso /mnt/iso iso9660 loop 0 0 mount -a
```

```
mount | grep examine
```

### Answer:

---

A

## Question 10

---

Question Type: MultipleChoice

Create a volume group, and set the size is 500M, the size of single PE is 16M. Create logical volume named lv0 in this volume group, set size is 20 PE, make it as ext3 file system, and mounted automatically under data.

### Options:

---

A- Explanation:

```
fdisk /dev/vda
```

```
pvcreate /dev/vda3
```

```
vgcreate --s 16M vg0 /dev/vda3
```

```
lvcreate --n lv0 --l 20 vg0
```

```
mkfs.ext3 /dev/mapper/vg0-lv0
```

```
mkdir /data
```

```
/etc/fstab:
```

```
/dev/mapper/vg0-lv0 /data ext3 defaults 0 0
```

```
mount --a
```

```
mount | grep data
```

### Answer:

---

A

## Question 11

---

Question Type: MultipleChoice

One Logical Volume is created named as myvol under vo volume group and is mounted. The Initial Size of that Logical Volume is 400MB. Make successfully that the size of Logical Volume 200MB without losing any dat

a. The size of logical volume 200MB to 210MB will be acceptable.

### Options:

---

A- Explanation:

First check the size of Logical Volume: `lvdisplay /dev/vo/myvol`

Make sure that the filesystem is in a consistent state before reducing:

```
# fsck -f /dev/vo/myvol
```

Now reduce the filesystem by 200MB.

```
# resize2fs /dev/vo/myvol 200M
```

It is now possible to reduce the logical volume. `#lvreduce /dev/vo/myvol -L 200M`

Verify the Size of Logical Volume: `lvdisplay /dev/vo/myvol`

Verify that the size comes in online or not: `df -h`

### Answer:

---

A

## Question 12

---

Question Type: MultipleChoice

---

Part 2 (on Node2 Server)

Task 3 [Managing Logical Volumes]

Create a new volume group in the name of datavg and physical volume extent is 16 MB

Create a new logical volume in the name of data1v with the size of 250 extents and file system must xfs

Then the logical volume should be mounted automatically mounted under /data at system boot time

### Options:

---

A- Explanation:

\*

```
[root@node2 ~]# lsblk
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT
vdb 252:16 0 5G 0 disk
vdb1 252:17 0 4.2G 0 part
vgrz-lvrz 253:2 0 4.1G 0 lvm /datarz
vdc 252:32 0 5G 0 disk
vdd 252:48 0 5G 0 disk
vde 252:64 0 10G 0 disk
[root@node2 ~]# parted /dev/vdc mklabel msdos
[root@node2 ~]# parted /dev/vdc mkpart primary 1MiB 4200MiB
[root@node2 ~]# parted /dev/vdc set 1 lvm on
*
[root@node2 ~]# udevadm settle
[root@node2 ~]# pvcreate /dev/vdc1
Physical volume '/dev/vdc1' successfully created.
[root@node2 ~]# vgcreate -s 16M datavg /dev/vdc1
Volume group 'datavg' successfully created
[root@node2 ~]# lvcreate -n datalv -L 4000M datavg
Logical volume 'datalv' created.
[root@node2 ~]# mkfs.xfs /dev/datavg/datalv
[root@node2 ~]# mkdir /data
[root@node2 ~]# blkid
/dev/mapper/datavg-datalv: UUID='7397a292-d67d-4632-941e-382e2bd922ce'
BLOCK_SIZE='512' TYPE='xfs'
*
[root@node2 ~]# vim /etc/fstab
UUID=7397a292-d67d-4632-941e-382e2bd922ce /data xfs defaults 0 0
[root@node2 ~]# mount UUID=7397a292-d67d-4632-941e-382e2bd922ce /data
[root@node2 ~]# reboot
[root@node2 ~]# df -hT
Filesystem Type Size Used Avail Use% Mounted on
/dev/mapper/datavg-datalv xfs 3.9G 61M 3.9G 2% /data
```

Answer:

A

To Get Premium Files for EX200 Visit

<https://www.p2pexams.com/products/ex200>

For More Free Questions Visit

<https://www.p2pexams.com/redhat/pdf/ex200>

**20%**  
**DISCOUNT**

**P2P**  
exams