Free Questions for EX200 by braindumpscollection

Shared by Hughes on 06-06-2022

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

Part 2 (on Node2 Server)

Task 8 [Tuning System Performance]

Set your server to use the recommended tuned profile

Options:

A- Explanation:
[root@node2 ~]# tuned-adm list
[root@node2 ~]# tuned-adm active
Current active profile: virtual-guest
[root@node2 ~]# tuned-adm recommend
virtual-guest
[root@node2 ~]# tuned-adm profile virtual-guest
[root@node2 ~]# tuned-adm active
Current active profile: virtual-guest
[root@node2 ~]# reboot
[root@node2 ~]# tuned-adm active
Current active profile: virtual-guest

Answer:

A

Question 2

Question Type: MultipleChoice

Part 2 (on Node2 Server)

Task 7 [Implementing Advanced Storage Features]

Create a thin-provisioned filesystem with the name think_fs from a pool think_pool using the devices.

The filesystem should be mounted on /strav and must be persistent across reboot

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
vdo1 253:4 0 50G 0 vdo /vbread
[root@node2 ~]# yum install stratis* -y
[root@node2 ~]# systemctl enable --now stratisd.service
[root@node2 ~]# systemctl start stratisd.service
[root@node2 ~]# systemctl status stratisd.service
[root@node2 ~]# stratis pool create think_pool /dev/vdd
[root@node2 ~]# stratis pool list
Name Total Physical Properties
think_pool 5 GiB / 37.63 MiB / 4.96 GiB ~Ca,~Cr
*
[root@node2 ~]# stratis filesystem create think_pool think_fs
[root@node2 ~]# stratis filesystem list
Pool Name Name Used Created Device UUID
think_pool think_fs 546 MiB Mar 23 2021 08:21 /stratis/think_pool/think_fs ade6fdaab06449109540c2f3fdb9417d
[root@node2 ~]# mkdir /strav
[root@node2 ~]# lsblk
[root@node2 ~]# blkid
/dev/mapper/stratis-1-91ab9faf36a540f49923321ba1c5e40d-thin-fs-ade6fdaab06449109540c2f3fdb9417d: UUID='ade6fdaa-b064-4910-9540-c2f3fdb9417d' BLOCK_SIZE='512' TYPE='xfs'
*
[root@node2 ~]# vim /etc/fstab
UUID=ade6fdaa-b064-4910-9540-c2f3fdb9417d /strav xfs defaults,x-systemd.requires=stratisd.service 0 0
[root@node2 ~]# mount /stratis/think_pool/think_fs /strav/
[root@node2 ~]# df -hT

/dev/mapper/stratis-1-91ab9faf36a540f49923321ba1c5e40d-thin-fs-ade6fdaab06449109540c2f3fdb9417d xfs 1.0T 7.2G 1017G 1% /strav

Answer:

A

Question 3

Question Type: **MultipleChoice**

Part 2 (on Node2 Server)

Task 6 [Implementing Advanced Storage Features]

Add a new disk to your virtual machine with a ize 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% /vbread

Answer:

A

Question 4

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 swapspace 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 5

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 6

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 datalv 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

Question 7

Question Type: MultipleChoice

Part 2 (on Node2 Server)

Task 2 [Installing and Updating Software Packages]

Configure your system to use this location as a default repository:

Also configure your GPG key to use this location

Options:

A- Explanation:

[root@node1 ~]# vim /etc/yum.repos.d/redhat.repo
[BaseOS]
name=BaseOS
baseurl=http://utility.domain15.example.com/BaseOS
enabled=1
gpgcheck=1
gpgkey=http://utility.domain15.example.com/RPM-GPG-KEY-redhat-release
[AppStream]
name=AppStream
baseurl=http://utility.domain15.example.com/AppStream
enabled=1
gpgcheck=1
gpgkey=http://utility.domain15.example.com/RPM-GPG-KEY-redhat-release
[root@node1 ~]# yum clean all
[root@node1 ~]# yum repolist
repo id repo name
AppStream AppStream
BaseOS BaseOS
[root@node1 ~]# yum list all

Answer:

A

Question 8

Question Type: MultipleChoice

Part 2 (on Node2 Server)

Task 1 [Controlling the Boot Process]

Interrupt the boot process and reset the root password. Change it to kexdrams to gain access to the system

Options:

A- Explanation:

*

1. Reboot the server pressing by Ctrl+Alt+Del
2. When the boot-loader menu appears, press the cursor keys to highlight the default boot-loader entry
3. Press e to edit the current entry.
4. Use the cursor keys to navigate to the line that starts with linux.
5. Press End to move the cursor to the end of the line.
6. Append rd.break to the end of the line.
7. Press Ctrl+x to boot using the modified configuration.
8. At the switch_root prompt

*

switch_root:/# mount -o remount,rw /sysroot

switch_root:/# chroot /sysroot

sh-4.4# echo kexdrams | passwd --stdin root
Changing password for user root.
passwd: all authentication tokens updated successfully.
sh-4.4# touch /.autorelabel
sh-4.4# exit; exit
*
Type exit twice to continue booting your system as usual.

Answer:

A

Question 9

Question Type: **MultipleChoice**

Part 1 (on Node1 Server)

Task 17 [Accessing Linux File Systems]

Find all the files owned by user "alex" and redirect the output to /home/alex/files.

Options:

A- Explanation:

* root@node1 ~]# find / -user alex -type f > /home/alex/files

Answer:

A

Question 10

Question Type: **MultipleChoice**

Part 1 (on Node1 Server)

Task 16 [Running Containers]

Configure your host journal to store all journal across reboot

Copy all journal files from /var/log/journal/ and put them in the /home/shangrila/container-logserver

Create and mount /home/shangrila/container-logserver as a persistent storage to the container as /var/log/ when container start

Options:

A- Explanation:

*

```
[shangrila@node1 ~]$ podman ps
CONTAINER ID IMAGE COMMAND CREATED STATUS PORTS NAMES
d5ffe018a53c registry.domain15.example.com:5000/rhel8/rsyslog:latest /bin/rsyslog.sh 5 seconds ago Up 4 seconds ago logserver
[shangrila@node1 ~]$ podman stats logserver
Error: stats is not supported in rootless mode without cgroups v2
[shangrila@node1 ~]$ podman stop logserver
d5ffe018a53ca7eb075bf560d1f30822ab6fe51eba58fd1a8f370eda79806496
[shangrila@node1 ~]$ podman rm logserver
Error: no container with name or ID logserver found: no such container
[shangrila@node1 ~]$ mkdir -p container-journal/
```

*

```
[shangrila@node1 ~]$ sudo systemctl restart systemd-journald
[sudo] password for shangrila:
[shangrila@node1 ~]$ sudo cp -av /var/log/journal/* container-journal/
[shangrila@node1 ~]$ sudo cp -av /var/log/journal/* container-journal/
[shangrila@node1 ~]$ sudo chown -R shangrila container-journal/
[shangrila@node1 ~]$ podman run -d --name logserver -v /home/shangrila/container-journal/:/var/log/journal:Z
registry.domain15.example.com:5000/rhel8/rsyslog
[shangrila@node1 ~]$ podman ps
[shangrila@node1 ~]$ loginctl enable-linger
[shangrila@node1 ~]$ loginctl show-user shangrila|grep -i linger
Linger=yes
```

*

```
[shangrila@node1 ~]$ podman stop logserver
```

[shangrila@node1 ~]$ podman rm logserver
[shangrila@node1 ~]$ systemctl --user daemon-reload
[shangrila@node1 ~]$ systemctl --user enable --now container-logserver
[shangrila@node1 ~]$ podman ps
CONTAINER ID IMAGE COMMAND CREATED STATUS PORTS NAMES
3903e1d09170 registry.domain15.example.com:5000/rhel8/rsyslog:latest /bin/rsyslog.sh 4 seconds ago Up 4 seconds ago logserver
[shangrila@node1 ~]$ systemctl --user stop container-logserver.service
*
[shangrila@node1 ~]$ sudo reboot
[shangrila@node1 ~]$ podman ps -a
CONTAINER ID IMAGE COMMAND CREATED STATUS PORTS NAMES
7e6cd59c506a registry.domain15.example.com:5000/rhel8/rsyslog:latest /bin/rsyslog.sh 10 seconds ago Up 9 seconds ago logserver

Answer:

A

Question 11

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 12

Question Type: MultipleChoice

Part 1 (on Node1 Server)

Task 14 [Managing SELinux Security]

You will configure a web server running on your system serving content using a non-standard port (82)

Options:

A- Explanation:

*

```
[root@node1 ~]# curl http://node1.domain15.example.com
curl: (7) Failed to connect to node1.domain15.example.com port 80: Connection refused
[root@node1 ~]# yum install httpd
[root@node1 ~]# systemctl enable --now httpd
Created symlink /etc/systemd/system/multi-user.target.wants/httpd.service /usr/lib/systemd/system/httpd.service.
[root@node1 ~]# systemctl start httpd
[root@node1 ~]# systemctl status httpd
Status: 'Running, listening on: port 80'
```

*

```
[root@node1 ~]# wget http://node1.domain15.example.com
2021-03-23 13:27:28 ERROR 403: Forbidden.
[root@node1 ~]# semanage port -l | grep http
http_port_t tcp 80, 81, 443, 488, 8008, 8009, 8443, 9000
[root@node1 ~]# semanage port -a -t http_port_t -p tcp 82
[root@node1 ~]# semanage port -l | grep http
http_port_t tcp 82, 80, 81, 443, 488, 8008, 8009, 8443, 9000
[root@node1 ~]# firewall-cmd --zone=public --list-all
[root@node1 ~]# firewall-cmd --permanent --zone=public --add-port=82/tcp
[root@node1 ~]# firewall-cmd --reload
[root@node1 ~]# curl http://node1.domain15.example.com
```

OK
*
root@node1 ~]# wget http://node1.domain15.example.com:82
Connection refused.
[root@node1 ~]# vim /etc/httpd/conf/httpd.conf
Listen 82
[root@node1 ~]# systemctl restart httpd
[root@node1 ~]# wget http://node1.domain15.example.com:82
2021-03-23 13:31:41 ERROR 403: Forbidden.
[root@node1 ~]# curl http://node1.domain15.example.com:82
OK

Answer:

A