(2½ Hours)

[Total Marks: 75]

N. B.: (1) **All** questions are **compulsory**.
(2) Make **suitable assumptions** wherever necessary and **state the assumptions** made.
(3) Answers to the **same question** must be **written together**.
(4) Numbers to the **right** indicate **marks**.
(5) Draw **neat labeled diagrams** wherever **necessary**.
(6) Use of **Non-programmable** calculators is **allowed**.

1. **Attempt *any three* of the following:** **15**

a. List and explain the duties of a system administrator.

Ans: **Creating and Maintaining User Accounts**

The system administrator creates a new user account, sets passwords for user account, decides whether to let users select their own passwords, assign passwords, allow users to change their passwords periodically. He can disable / delete the user account, depending on how crucial / essential is the data contained in the account
Examples of creation of user accounts and passwords and policies.

b. Write a short note on bash.

Ans: **Bash Shell:**

To communicate commands to the operating system kernel, an interface is needed that sits between the kernel and the end user issuing these commands. This interface is known as the shell. Several shells are available on RHEL. Bash (Bourne Again Shell) is the one that is used in most situations. This is because it is compatible with the Bourne shell, which is commonly found on UNIX servers

Features of Bash:

Autocompletion: Bash offers the option to complete a command automatically. Bash can complete almost everything, not just commands. It can also complete filenames and shell variables. To use this nice feature of completion, use the Tab key. An example of how this works follows. In this example, the cat command is used to display the contents of an ASCII text file. The name of this file, which is in the current directory, is this_is_a_file. To open this file, the user can type cat thi and then immediately hit the Tab key. If there is just one file that starts with the letters thi, Bash will automatically complete the name of the file. If there are more options, Bash will complete the name of the file as far as possible. This happens, for example, when in the current directory there is a file with the name this_is_a_text_file and thisAlsoIsAFile. Since both files start with this, Bash completes only up to this and doesn't go any further. To display a list of possibilities, you can then hit the Tab key again. This allows you to enter more information manually. Of course, you can then use the Tab key to use the completion feature again.

History: The history mechanism helps you remember the last commands you used. By default, the last 1,000 commands of any user are remembered. History allows you to use the up and down arrow keys to navigate through the list of commands that you used previously. You can see an overview of these remembered commands when using the history command from the Bash command line. This command shows a list of all of the recently used commands. From this list, a command can also be restarted. For example, if you see command 5 in the list of commands, you can easily rerun this command by using its number preceded by an exclamation mark, or !5 in this example. The file .bash_history, stores all of the commands you have used before. history -c will clear the history list for the user who uses this command.

Bash key sequences:

Ctrl+C Use this key sequence to quit a command that is not responding (or simply is taking

too long to complete).
Ctrl+D This key sequence is used to send the end-of-fi le (EOF) signal to a command.
Ctrl+R This is the reverse search feature. When used, it will open the reverse-i-search prompt. This feature helps you locate commands you have used previously.
Ctrl+Z It stops your command, but it does not terminate it.
Ctrl+A This keystroke brings the cursor to the beginning of the current command line.
Ctrl+B This keystroke moves the cursor to the end of the current command line.

c. Discuss the difference between piping and redirection with the help of an example.

Ans: **Piping:** It is used to send the result of a command to another command

• Open a shell, and use su - to become the root. Enter the root password when prompted.
• Type the command ps aux. This command provides a list of all the processes that are currently running on your computer. You'll notice that the list doesn't fit on the screen.
• To make sure you can see the complete result page by page, use ps aux | less. The output of ps is now sent to less, which outputs it so that you can browse it page by page.

Redirection: It sends the output of a command to a file.

• From a console w indow, use the command ps aux. You'll see the output of the command on the current console.
• Now use ps aux > ~/psoutput.txt. You don't see the actual output of the command, because it is w rit ten to a fi le that is created in your home directory, which is designated by the ~ sign.
• To show the contents of the file, use the command less ~/psoutput.txt.

d. With the help of an example, write the steps to mount a device in Linux.

Ans: Eg: To mount a USB device, follow the below mentioned steps:

1. Open a terminal, and make sure you have root privileges.
2. Insert a USB flash drive in the USB por t of your computer.
3. Use dmesg to find the device name of the USB fl ash drive. (asume it is /dev/sdb)
4. Use fdisk -cul /dev/sdb to fi nd current par titions on the USB fl ash drive. I'll assume you will find one partition with the name of /dev/sdb1.
5. Use mount /dev/sdb1 /mnt to mount the USB flash drive on the /mnt directory.
6. Use cd /mnt to go into the /mnt directory.
7. Type ls to verify that you see the contents of the USB flash drive.
8. Use cd without any arguments. This takes your current shell out of the /mnt directory and back to your home directory.
9. Use umount /dev/sdb1 to try to dismount the USB flash drive.

e. List the differences between hard and symbolic links.

Ans: **Symbolic Link:**

A symbolic link is the most flexible link type you can use. It points to any other file and any other directory, no matter where it is. With symbolic links, there is a difference between the original file and the link. If you remove the original file, the symbolic link won't work anymore and thus is invalid.

Hard link:

A hard link can be used only to point to a file that exists on the same device. A hard link is more like an additional name you'd give to a file. The inode is the administration of a file. To get to a file, the file system reads the file's inode in the file system metadata, and from there it learns how to access the block where the actual data of the file is stored. To get to the inode, the file system uses the filename that exists somewhere in a directory. A hard link is an additional filename that you can create anywhere in a directory on the same device that gives access to the same file system metadata. With hard links, you only need the original filename to create the hard link. Once it has been created, it isn't needed anymore, and the original filename can be removed.

f. Explain the steps to create and manage your own repository in yum.

Ans: **Creating Repositories:**

1. Use mkdir /repo to create a directory that you can use as a repository in the root of your server's file system.
2. Insert the Red Hat installation DVD in the optical drive of your server. Assuming that you run the server in graphical mode, the DVD will be mounted automatically.
3. Use the cd /media/RHEL[Tab] command to go into the mounted DVD. Next use cd Packages, which brings you to the directory where all RPMs are by default. Now use cp * /repo to copy all of them to the /repo directory you just created. Once this is finished, you don't need the DVD anymore.
4. Now use cd /repo to go to the /repo directory. From this directory, type rpm -ivh createrepo[Tab]. This doesn't work, and it gives you a " Failed dependencies" error. To install createrepo, you fi rst need to install the deltarpm and python-deltarpm packages. Use rpm -ivh deltarpm[Tab] python-deltarpm[Tab] to install both of them. Next, use rpm -ivh createrepo[Tab] again to install the createrepo package.
5. Once the createrepo package has been installed, use createrepo /repo, which creates the metadata that allows you to use the /repo directory as a repository. This will take a few minutes. When this procedure is finished, your repository is ready for use.

Managing Repositories:

1. Use the command yum repolist. In its output (repolist: 0), the command tells you that currently no repositories are configured.
2. Use the command yum search nmap. The result of this command is the message No Matches found.
3. Now use vi to create a file w ith the name /etc/yum.repos.d/myrepo.repo. Note that it is impor tant that the fi le has the extension .repo. The file should have the following contents:

```
[myrepo]
name=myrepo
baseurl=file:///repo
gpgcheck=0
```

4. Now use the commands yum repolist and yum search nmap again.

2. Attempt *any three* of the following: 15

a. Explain the different kinds of partitions in linux and their characteristics.

Ans: There are three kinds of partitions available in linux:

Primary Partitions: These are written directly to the master boot record of your hard drive. After creating four primary partitions, you can't add any more partitions—even if there is still a lot of disk space available. There's space for just four partitions in the partition table and no more than four.

Extended Partition: Every hard drive can have one extended partition. You cannot create a file system in an extended partition. The only thing you can do with it is to create logical partitions. You'll use an extended partition if you intend to use more than four partitions in total on a hard drive.

Logical Partitions: A logical partition is created inside an extended partition. You can have a maximum of 11 logical partitions per disk, and you can create fi le systems on top of logical partitions.

b. Discuss the steps to create a swap file.

Ans: Swap space is allocated when your server is completely out of memory and using swap space allows your server to continue to offer its services. At least a minimal amount of swap space should always be available.

Steps to create swap file:

1. Use dd if=/dev/zero of=/swapfile bs=1M count=1024. This command creates a 1GB swap file in the root directory of your server.
2. Use mkswap /swapfile to mark this file as swap space.
3. Type free -m to verify the current amount of swap space on your server. This amount is expressed in megabytes.

4. Type swapon /swapfile to activate the swap file.
5. Type free -m again to verify that you just added 1GB of swap space.
6. Open /etc/fstab with an editor, and put in the following line:
/swapfile swap swap defaults 0 0

c. What are runlevels in Linux? Explain the commands used to manage services.

Ans: **Runlevels in Linux:**

The runlevel defines the state in which the server boots. Every runlevel is referenced by number. Common runlevels are runlevel 3 and runlevel 5. Runlevel 3 is used to start services that are needed on a server that starts without a graphical user interface, and runlevel 5 is used to define a mode where the runlevel starts with a graphical interface. In each runlevel, service scripts are started. These service scripts are installed in the /etc/init.d directory and managed with the service command.

Commands to manage services:

To manage service scripts, two commands are relevant – service and chkconfig

service command:

Used to start, stop, and monitor all of the service scripts in the /etc/init.d directory.
Eg: service dhcpd start – will start the dhcpd service if it is not running.
service dhcpd stop – will stop the dhcpd service if it is running.
service dhcpd restart – will stop and then start the dhcpd service.
service dhcpd status – will show the current status of the dhcpd service whether is is running or stopped

chkconfig command:

Used to enable the service in the runlevel
For eg: You may have started ntpd, but after a reboot it won't be started automatically. Use chkconfig ntpd on to add the ntpd service to the runlevels of your server. To verify that ntpd has indeed been added to your server 's runlevels, type chkconfig –list. This command lists all services and their current status. You can filter the results by adding grep ntpd to the chkconfig --list command.

d. Discuss the steps to configure key based SSH authentication.

Ans: **Key based SSH authentication:**

The default authentication method in SSH is password based, which means when connecting to a server, you need to enter the password of the user with whom you are connecting. But someone can guess your password and if you frequently need to connect to the same server, it's a waste of time to enter the identical password over and over.

Alternative is to use key based authentication. When SSH key-based authentication is used, make sure the public key is available on the servers to all users who need to use this technology where they want to log in. When logging in, the user creates an authentication request that is signed with their private key. This authentication request is matched to the public key of the same user on the server where that user wants to authenticate. If it matches, the user is allowed to enter; if it doesn't, the user is denied access.

Steps to configure key based SSH authentication:

1. Open a root shell, and from there use ssh server. You will be prompted for a password.
2. Type exit to close the SSH session.
3. Now generate the public-private key pair using ssh-keygen. You will be prompted for the file in which you want to save the private key. Press Enter to accept the default, which saves the key in /root/.ssh/id_rsa. Next press Enter twice to save the key w ithout a passphrase. This completes the procedure and creates two fi les: id_rsa and id_rsa.pub. The private key is stored in id_rsa, and the public key is stored in id_rsa.pub.
4. Now you need to copy the public key to the server where you want to use it. Use ssh_copy_id server to do this. This copies the public key to the server and generates some messages.

5. Use ssh server to connect to the server again. You'll notice that you won't be prompted for a password because the SSH keys are used to establish the connection.

e. Elaborate what basic permissions are and how they are applied to files and directories in linux.

Ans:

Basic Permission	Applied to files	Applied to directories
Read	Open a file	List contents of a directory
Write	Change contents of a file	Create and delete files
Execute	Run a program file	Change to the directory

The three basic permissions allow you to read, write, and execute files. The effect of these permissions is different if applied to files vs. directories.

Read permission: If applied to a file, the read permission gives you the right to open the file for reading. If applied to a directory, read allows you to list the contents of that directory. But this permission does not allow you to read files in the directory.

Write permission: It allows you to modify the contents of existing files. It does not, allow you to create or delete new files. To do that, you need write permission on the directory where you want to create the fi le. In directories, this permission also allows you to create and remove new subdirectories and fi les, but you need execute as well to descend into the directory.

Execute permission: It is required to execute a file. But it is never set by default. While the execute permission on files means you are allowed to run a program file, when applied to a directory, it indicates that the user can use the cd command to go to that directory.

f. Explain the user information configuration file.

Ans: /etc/passwd is the user information configuration file.

Different fields are used in /etc/passwd.

Username The user's login name is stored in the first field in /etc/passwd.

Password In the old days of UNIX, encrypted passwords were stored in this file.

UID As you have already learned, every user has a unique user ID. Red Hat Enterprise Linux starts numbering local user IDs at 500, and typically the highest number that is used is 60000 (the highest numbers are reserved for special-purpose accounts).

GID Every user has a primary group. The group ID of this primary group is listed there. On Red Hat Enterprise Linux, every user is also a member of a private group that has the name of the user.

GECOS The General Electric Comprehensive Operating System (GECOS) field is used to include some additional information about the user. The field can contain anything you like, such as the department where the user works, the user's phone number, or anything else. This makes identifying a user easier for an administrator. The GECOS field is optional.

Home Directory This field points to the directory of the user's home directory.

Shell The last field in /etc/passwd is used to refer to the program that is started automatically when a user logs in. Most often, this will be /bin/bash

Eg: for user linda the entry in the file will look like

linda:x:500:500:linda:/home/linda:/bin/bash

3. Attempt *any three* of the following: **15**

a. State the steps to setup a firewall that allows SSH packets.

Ans: 1. On the host computer, type iptables -L -v to display the current configuration.

2. Type the following commands: iptables -P INPUT ACCEPT, iptables -P OUTPUT ACCEPT, and iptables -P FORWARD ACCEPT. Now use iptables -F to flush all other rules.

3. Use iptables -L -v to verify that the policy is set to ACCEPT for all three chains in the filter table. Use service iptables save

4. Repeat steps 1 to 3 on the virtual machine.

5. After clearing the firewall on the virtual machine, you must test which services are offered from the virtual machine. Use ping to test whether you can still reach the virtual machine
6. Next use yum install -y nmap on the host to install the nmap network scanner.
7. After installing it, use nmap, followed by the IP address of the vir tual machine to scan available services on the virtual machine
8. Set a policy for the three chains. To do this, type the following:
   iptables -P INPUT DROP
   iptables -P OUTPUT DROP
   iptables -P FORWARD DROP
9. Open the loopback interface first. Type
   iptables -A INPUT -i lo -j ACCEPT
   iptables -A OUTPUT -o lo -j ACCEPT
10. Open the SSH port. Type
    iptables -A INPUT -p tcp --dport 22 -j ACCEPT
    iptables -A OUTPUT -m state --state ESTABLISHED,RELATED -j ACCEPT

b. What are modules in a firewall? Explain the limit module.

Ans: **Modules:**

A module is an optional element that you can use in a rule. Modules offer enhancements to the Netfilter firewall. They do that by loading a specific kernel module that adds functionality.

Limit module:

When using logging, numerous lines are written to the log files. On an active system, the number of lines logged can be so high that your log files will suffer a denial-of-service attack, and nothing further will be readable. To protect yourself against this, you can use the limit module. This module allows you to jump to a target if a certain number of packets have been matched within a specific period of time.

Eg: if you want to limit the number of packets logged to one per second for all SSH traffic, use the following line:

iptables -I INPUT 3 -p tcp --dport 22 -m limit --limit 1/s -j LOG

Another way the limit module can be used to log incoming packets is by including it after all rules that allow incoming traffic. To know which traffic is denied by your firewall, type

iptables -A INPUT -m limit --limit 15/minute -j LOG

iptables -A OUTPUT -m limit --limit 15/minute -j LOG

c. Explain how to create and manage certificates with openssl.

Ans: Steps to create a self-signed certificate using openssl

1. To create a certificate for your CA server, use the configuration file in /etc/pki/openssl.cnf
2. After checking the default values you want to use in openssl.cnf, you can start creating your own self-signed certificate. The following command allows you to create a certificate that uses a 1024-bit RSA key with a validity of 10 years:
   openssl req -newkey rsa:1024 -x509 -days 3650

A more complex example is as follows, where a specific mention is made of where the private key and certifi cate must be written.

openssl req -newkey rsa:2048 -x509 -days 3650 -keyout private/my-CAkey.pem -out my-CAcert.pem

When you are creating a key, the openssl command prompts for a passphrase as it cannot be read from a configuration file.

d. List the steps to encrypt, share and decrypt files using GPG.

Ans: Steps to encrypt and decrypt files using gpg:

1. Open a shell and use su - linda to become user linda.
2. As linda, copy the fi le /etc/hosts to your home directory using cp /etc/hosts ~.
3. Use gpg --listkeys to list the keys currently impor ted in Linda's environment, and note the exact name of the user lisa.

4. Encrypt the file using gpg -e hosts. When the user account is requested, enter the exact name of user lisa as you found it in the previous step of this exercise. Next press Enter on an empty line to complete the encryption procedure.
5. Use cp ~/hosts.gpg /tmp to copy the gpg file to the tmp directory where lisa can see and read it.
6. Use exit to log out as linda, and now use su - lisa to become user lisa.
7. As lisa, use gpg -d /tmp/hosts.gpg to decrypt the hosts file.

e. With the help of an example, explain exporting and mounting of NFS share.

Ans: On NFS server, perform the following steps:
1. Use mkdir /data to create a directory with the name /data.
2. After creating the directory you want to share, make sure that incoming users also have permissions on that share. To do this, you'll use regular Linux permissions. In this case, use chmod 777 /data to open the share to anyone.
3. Use an editor to open the file /etc/exports, and insert the following line of code to share the /data directory with anyone:
/data *(rw, all_squash)
4. Use service nfs restart to restart the NFS server and have it offer the new share.
5. Use chkconfig nfs on to make sure that the NFS server starts at a server reboot.
6. To verify that the share is available, use showmount -e localhost, which shows all exported shares on the local computer.

On NFS client give the following command as root
Mount -o nolock *ipaddressOfNFSServer*:/data /data

f. Discuss the steps to setup a Samba server.

Ans: Samba is a very versatile service that you can use for different purposes on your network. Apart from sharing fi les, it can share printers and also offer Windows domain services such as directory services. Samba can be integrated into an Active Directory domain and made a member server of Active Directory.
Steps to setup samba server:
1. Create a directory on the Linux file system on the Samba server.
2. If needed, create Linux users and give the appropriate permissions to the directory you just created.
3. Install the Samba server.
4. Define the share in /etc/samba/smb.conf.
5. Create a Samba user account that has access to the share.
6. (Re)start the Samba service.
7. Tell SELinux to give access to the Samba share.

4. Attempt *<u>any three</u>* of the following: **15**

a Write a short note on cache-only nameserver.

Ans: Open a terminal, log in as root, and run yum -y install bind-chroot on the host computer to install the bind package.
2. With an editor, open the configuration fi le /etc/named.conf. You need to change some parameters in the configuration file to have BIND offer its services to external hosts.
3. Change the file to include the following parameters: listen-on port 53 { any; }; and allow-query { any; };. This opens your DNS server to accept queries on any network interface from any client.
4. Still in /etc/named.conf, change the parameter dnssec-validation; to dnsserver-validation no;.
5. Finally, insert the line forwarders x.x.x.x in the same configuration file, and give it the value of the IP address of the DNS server you normally use for your Internet connection. This ensures that the DNS server of your Internet provider is used for

DNS recursion and that requests are not sent directly to the name servers of the root domain.
6. Use the service named restart command to restart the DNS server.
7. From the RHEL host, use dig redhat.com. You should get an answer, which is sent by your DNS server. You can see this in the SERVER line in the dig response.

b Explain the DHCP server configuration.

Ans: The Dynamic Host Configuration Protocol (DHCP) is used to assign IP-related configuration to hosts in your network.
Steps to setup a DHCP server:
1. Start the virtual machine, and open a root shell. From the root shell, use the command
yum -y dhcp to install the DHCP server.
2. Open the file /etc/dhcp/dhcpd.conf with an editor, and give it the following contents. Make sure that the names and IP addresses used in this example match your network:

```
option domain-name "example.com";
option domain-name-servers YOUR.DNS.SERVERNAME.HERE;
default-lease-time 600;
max-lease-time 1800;
subnet 192.168.100.0 netmask 255.255.255.0 {
range 192.168.100.10 192.168.100.20;
options routers 192.168.100.1;
}
```

3. Start the DHCP server by using the command service dhcpd start, and enable it using chkconfig dhcpd on.
4. Start the second virtual machine. Make sure that the network card is set to get an IP address from a DHCP server. After star ting it, verify that the DHCP server has indeed handed out an IP address.

c Discuss the role of MUA, MTA, and MDA in the email process.

Ans: **Message Transfer Agent (MTA):** The MTA uses the Simple Mail Transfer Protocol (SMTP) to exchange mail messages with other MTAs on the Internet. If a user sends a mail message to a user on another domain on the Internet, it's the responsibility of the MTA to contact the MTA of the other domain and deliver the message there. To find out which MTA serves the other domain, the DNS MX record is used. Upon receiving a message, the MTA checks whether it is the final destination. If it is, it will deliver the message to the local message delivery agent (MDA), which takes care of delivering the message to the mailbox of the user. If the MTA itself is not the final destination, the MTA relays the message to the MTA of the final destination. MTA relay messages only for authenticated users or users who are known in some other way. If, for some reason, the MTA cannot deliver the message to the other MTA, it will queue it. Upon delivery, it sometimes happens that the MTA, which contacted an exterior MTA and delivered the message there, receives it back.
Mail Delivery Agent (MDA): Upon receiving a message, the MTA typically delivers it at the mail delivery agent. The MDA delivers mail to the recipient's local message store, /var/spool/mail/$USER. In the Postfix mail server, an MDA is included in the form of the local program. MDA is only the software part that drops the message somewhere the recipient can find it.
Message User Agent (MUA): The mail message arrives in the mail user agent (MUA). This is the mail client that end users use to read their messages or to compose new messages. It is the responsibility of users to install an MUA, which allows them to work with email on their computer, tablet, or smartphone. Popular MUAa are Outlook, Evolution, and the Linux command-line Mutt tool

d Explain the various parameters for secure internet configuration of Postfix server.

Ans: To make a secure Internet configuration, you need to set the following parameters in the /etc/postfix/main.cf file.

myhostname This parameter specifies the name of this host. If not specified, it is set to the full DNS domain name (FQDN) of this host. This parameter is used as a variable in other parameters in the main.cf file, so it is useful to set it.

mydomain This parameter specifi es the domain of this host. If not set, the domain name part of the FQDN is used.

myorigin This parameter determines the domain seen by the email recipient when receiving messages. The default is to use the FQDN of this host. This means that if user linda on server dfw.example.com sends a message, the recipient will see a message coming in from linda@dfw.example.com. To append the domain name only and not the entire FQDN, use myorigin = $mydomain.

inet_interfaces This parameter specifies the IP addresses of the mail server to which it binds. By default, it is set to localhost only, which means that your mail server cannot receive messages from the Internet. To enable all inet_interfaces use inet_interfaces = all.

mynetworks This parameter is optional. You can use it to specify the network address from which your MTA accepts messages for relaying without further authentication.

e State the steps to setup virtual hosts in Apache.

Ans: To handle more than one site from an Apache server, you can create virtual hosts. A virtual host is a definition of different websites to be served by the Apache web server. This definition can be included in the main Apache configuration file /etc/httpd/conf/httpd.conf or in separate files that can be created in the /etc/httpd/conf.d/ directory.

Either a name-based virtual host, an IP-based virtual host, or both can be configured. Name-based virtual hosts are the default, and they are easier to set up because multiple Apache sites can be run on one IP address. IP-virtual hosts are often used if SSL is needed on a website, because in SSL it is beneficial if a connection can be traced back to its original unique IP address. When setting up a virtual host, almost any Apache directive can be used.

Steps to setup Virtual Hosts:

1. On your host computer, open the file /etc/hosts with an editor and add two lines that make it possible to resolve the names of the virtual host you are going to create to the IP address of the virtual machine.
2. On the virtual machine, open a root shell and create a configuration file with the name server1.example.com.conf in the directory /etc/httpd/conf.d. Give this file the following content :

<VirtualHost *:80>
ServerAdmin webmaster@server1.example.com
DocumentRoot /www/docs/server1.example.com
ServerName server1.example.com
ErrorLog logs/server1/example.com-error_log
CustomLog logs/server1.example.com-access_log common
</VirtualHost>

3. Close the confi guration file, and from the root shell, use mkdir -p /www/docs/ server1.example.com.
4. Create a file w ith the name index.html in the server1 document root, and make sure its contents read “ Welcome to server1.”
5. Use semanage fcontext -a -t httpd_sys_content_t "/www(/.*)", followed by restorecon -r /www.
6. Open the /etc/httpd/conf/httpd.conf file with your editor, and take out the # sign at the front of the line NameVirtualHost *:80.

7. Use service httpd restart to restart the Apache web server.
8. Use elinks http://server1.example.com. You should now see the server1 welcome page.
9. Restart httpd, and verify that server1 is accessible.

f Explain how the DirectoryIndex, Options, AllowOverride, and Order directives in Apache.

Ans: **DirectoryIndex:** This directive can be used to specify that other files should also be considered. If this is the case, it will show the contents of this file, and if not, a list of files in the directory is shown.
To modify this behavior, the DirectoryIndex and Options directives can be used. By default, the DirectoryIndex directive specifies that Apache should look for a file with the name index.html or index.html.var.

Options: This directive within a directory definition can further fine-tune the options that are used to display the contents of a directory. Options can also be used to determine which server features are available in a particular directory. A useful argument for the Options directive is Indexes. This option will show a list of files in the directory if no index.html is available. Related to this option is FollowSymLinks. This option will ensure that symbolic links are followed if they exist in the document directory.

AllowOverride: This directive is related to the .htaccess file that an administrator can use to restrict access to a given directory. If AllowOverride is set to none, the contents of any .htaccess file that is found anywhere in a subdirectory of the current directory will be ignored. If you don't want the owners of subdirectories to restrict access to their directories, set AllowOverride to none. If you want to allow users to restrict access to subdirectories, set it to All.

Order: Allows to handle access restrictions in a basic way. With this directive, the order in which allow and deny commands are used, can be specified. The order is not defined by how the rules appear in your configuration fi le but by how you've used the Order directive. The default order is deny and then allow. This means that if a client is excluded by deny, it will be excluded unless it matches allow. If neither is matched, the client gets access.
Eg:
order deny, allow
allow from 10.100
deny from all
The httpd process first reads the deny line, which denies access to all, and then handles all exceptions as stated on the allow lines.

5. Attempt *any three* of the following: **15**

a. Discuss the various ways in which a shell script can be executed.

Ans: A shell script is a text file that contains a sequence of commands. Basically, anything that can run a bunch of commands is considered a shell script.
There are three different ways of executing a shell script
• Make it executable, and run it as a program.
After making the script executable using a command like chmod +x hello, you can run it just like any other command. The only limitation is the exact location in the directory structure of your script. If it is in the search path, you can run it by typing any command. If it is not in the search path, you have to run it from the exact directory where it is located. This means that if user linda created a script with the name hello in /home/linda, she has to run it using the command /home/linda/hello. Alternatively, if she is already in /home/linda, she could use ./hello to run the script.

In the latter example, the dot and slash tell the shell to run the command fromthe current directory.

• Run it as an argument of the bash command.

The second option for running a script is to specify its name as the argument of the bash command. For example, the script hello would run using the command bash hello. The advantage of running the script this way is that there is no need to make it executable first. If you run it this way, you can specify an argument to the bash command while running it.

Make sure you are using a complete path to the location of the script when running it this way. It has to be in the current directory, or you would have to use a complete reference

to the directory where it is located. This means that if the script is /home/linda/hello and

your current directory is /tmp, you should run it using bash /home/linda/hello

• Source it.

The third way of running a script is completely different. You can source the script. By sourcing a script, you don't run it as a subshell. Rather, you include it in the current shell. This can be useful if the script contains variables that you want to be active in the current shell. There are two ways to source a script. These two lines show you how to source a script that has the name settings:

```
. settings
source settings
```

b. Write a shell script to add ten users, remove their passwords and add them to the group students. Use for loop.

Ans:
```
#!/bin/bash
groupadd students
for ((i=1;i<=10;i++)); do
    useradd user$i
    passwd -d user$i
    usermod -aG students user$i
done
```

c. Explain the steps to setup a quorum disk.

Ans: 1. On one cluster node, use fdisk to create a par tition on the iSCSI device. It doesn't need to be big—100MB is sufficient.

2. On the other cluster node, use the partx -a command to update the partition table. Now check /proc/partitions on both nodes to verify that the partition on the iSCSI disk has been created.

3. On one of the nodes, use the following command to create the quorum disk: mkqdisk

-c /dev/sdb1 -l quorumdisk. Before typing this command, make sure to doublecheck the name of the device you are using.

4. On the other node, use mkqdisk -L to show all quorum disks. You should see the quorum disk with the label quorumdisk that you just created.

5. In Conga, open the Configuration QDisk tab. On this tab, select the option Use A Quorum Disk

6. use the cman_tool status command to verify that it works as expected

d. Discuss the steps to setup fencing.

Ans: Fencing is needed to maintain the integrity of the cluster. Hardware fencing means that a hardware device is used to terminate a failing node. Typically, a power switch or integrated management card, such as HP ILO or Dell Drac, is used for this purpose.

To set up fencing, you need to perform two different steps.

First you need to configure the fence devices, and then you associate the fence devices to the nodes in the network. To define the fence device, you open the

Fence Devices tab in the Conga management interface. After clicking Add, you'll see a list of all available fence devices. A popular fence device type is IPMI LAN. This fence device can send instructions to many integrated management cards, including the HP ILO and Dell Drac. After selecting the fence device, you need to define its properties. These properties are different for each fence device, but they commonly include a username, a password, and an IP address. After entering these parameters, you can submit the device to the configuration.

e. Explain the steps to perform an automated installation using a kickstart file.

Ans: 1. On the installation server, copy the anaconda-ks.cfg fi le from the /root directory to the /www/docs/server1.example.com directory. You can just copy it straight to the root directory of the Apache virtual host. After copying the file, set the permissions to mode 644, or else the Apache user will not be able to read it.

2. Start Virtual Machine Manager, and click the Create Virtual Machine button. Enter a name for the virtual machine, and select Network Install.

3. On the second screen of the Create A New Virtual Machine Wizard, enter the URL to the web server installation directory: http://server1.example.com/install.
Open the URL options, and enter this Kickstar t URL: http://server1.example .com/anaconda-ks.cfg.

4. Accept all the default options in the remaining windows of the Create A New Virtual Machine Wizard, which will star t the installation. In the beginning of the procedure, you'll see the message Retrieving anaconda-ks.cfg. If this message disappears and you don't see any error messages, this indicates that the kickstar t file has loaded correctly.

f. List the steps to configure DHCP for PXE boot.

Ans: Add the following lines to dhcpd.conf

```
option space pxelinux;
option pxelinux.magic code 208 = string;
option pxelinux.configfile code 209 = text;
option pxelinux.pathprefix code 210 = text;
option pxelinux.reboottime code 211 = unsigned integer 32 ;
subnet 192.168.1.0 netmask 255.255.255.0 {
option routers 192.168.1.1 ;
range 192.168.1.200 192.168.1.250 ;
class "pxeclients" {
match if substring (option vendor-class-identifier, 0, 9) =
"PXEClient";
next-server 192.168.1.70;
filename "pxelinux/pxelinux.0";
}
}
```
