

- 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 two of the following:

10

a List and explain any five duties of System Administrator.

- Installing and Configuring Servers
- Installing and Configuring Application Software
- Creating and Maintaining User Accounts
- Backing Up and Restoring Files
- Monitoring and Tuning Performance
- Configuring a Secure System
- Using Tools to Monitor Security

b Explain the linux booting process.

When you turn on your PC, it runs a program called the basic input/output system (BIOS). the system BIOS is typically installed on a read-only memory (ROM) chip physically attached to the system board.

After the BIOS loads, it performs some diagnostics on the hardware, checks the installed components to be sure they are functioning, and checks the system RAM. Next, the BIOS tries to find a system drive from which it can load the boot program to begin the process of starting the operating system.

The first sector of the drive has an area called the Master Boot Record (MBR), which holds the program that is used to begin the actual loading of the operating system. As soon as the BIOS finds the MBR, it gives up control of the boot process. In the case of Fedora Core and Enterprise Linux, a program called a boot loader begins the loading of the operating system. The boot loader program used is called the *Grand Unified Boot loader*, or GRUB.

c What is the purpose of grub.conf file in linux? [1 mark]

Write the purpose of the following commands used in grub.conf file [4 mark]

The /boot/grub/grub.conf file controls what information is displayed on the graphical screen. This file even controls whether you see the graphical screen at all.

1. hiddenmenu : This command tells GRUB not to display the menu and to boot the default after the timeout expires.
2. title : This command tells GRUB to list a boot name on the menu using the name following the title command.
3. default : This command tells GRUB to boot the first listing beginning with title.
4. Timeout : This command tells GRUB to boot the default entry after five seconds.

d What is RAID? [2 mark]

Write different levels of RAID. [3 mark]

RAID is an acronym for Redundant Array of Inexpensive, or Independent (depending on who you ask), Disks. There are two types of RAID that can be used on computer systems. These types are hardware RAID and software RAID.

In **hardware RAID** the disks have their own RAID controller with built-in software that handles the RAID disk setup, and I/O. The controller is typically a card in one of the system's expansion

slots, or it may be built onto the system board. The hard RAID interface is transparent to Linux, so the hardware RAID disk array looks like one giant disk. The operating system does not control the RAID level used, it is controlled by the hardware RAID controller. Most dedicated servers use a hardware RAID controller.

In **software RAID** there is no RAID controller card. The operating system is used to set up a logical array, and the operating system controls the RAID level used by the system.

RAID Levels :

RAID level 0 — This RAID level requires at least two disks and uses a method called striping that writes data across both drives. There is no redundancy provided by this level of RAID, since the loss of either drive makes it impossible to recover the data. This level of RAID does give a speed increase in writing to the disks.

RAID level 1 — This RAID level requires at least two disks and uses a method called mirroring. With mirroring, the data is written to both of the drives. So, each drive is an exact mirror of the other one, and if one fails the other still holds all the data. There are two variants to level 1

with one variant using a single disk controller that writes to both disks as described above. The other variant uses two disk controllers, one for each disk. This variant of RAID level 1 is known as duplexing.

RAID level 5 — This RAID level, which is the most widely used, requires at least three disks and uses striping to write the data across the two disks similarly to RAID level 1. But unlike RAID level 1, this level of RAID uses the third disk to hold parity information that can be used to reconstruct the data from either, but not both, of the two disks after a single disk failure.

2 Attempt any two of the following:

10

a Explain any five NFS export options. [1 mark each]

| OPTION | DESCRIPTION |
|-----------------------------|---|
| <code>all_squash</code> | Maps all requests from all UIDs or GIDs to the UID or GID, respectively, of the anonymous user. |
| <code>anongid=gid</code> | Sets the GID of the anonymous account to <code>gid</code> . |
| <code>anonuid=uid</code> | Sets the UID of the anonymous account to <code>uid</code> . |
| <code>async</code> | Allows the server to cache disk writes to improve performance. |
| <code>fsid=n</code> | Forces NFS's internal file system identification (FSID) number to be <code>n</code> . |
| <code>hide</code> | Hides an exported file system that is a subdirectory of another exported file system. |
| <code>insecure</code> | Permits client requests to originate from unprivileged ports (those numbered 1024 and higher). |
| <code>insecure_locks</code> | Disables the need for authentication before activating lock operations (synonym for <code>no_auth_nlm</code>). |
| <code>mp [=path]</code> | Exports the file system specified by <code>path</code> only if the corresponding mount point is mounted (synonym for <code>mountpoint [=path]</code>). |
| <code>no_all_squash</code> | Disables <code>all_squash</code> . |
| <code>no_root_squash</code> | Disables <code>root_squash</code> . |

| | |
|-------------------------------|---|
| <code>no_subtree_check</code> | Disables <code>subtree_check</code> . |
| <code>no_wdelay</code> | Disables <code>wdelay</code> (must be used with the <code>sync</code> option). |
| <code>nohide</code> | Does not hide an exported file system that is a subdirectory of another exported file system. |
| <code>ro</code> | Exports the file system read-only, disabling any operation that changes the file system. |
| <code>root_squash</code> | Maps all requests from a user ID (UID) or group ID (GID) of 0 to the UID or GID, respectively, of the anonymous user (-2 in Red Hat Linux). |
| <code>rw</code> | Exports the file system read-write, permitting operations that change the file system. |
| <code>secure</code> | Requires client requests to originate from a secure (privileged) port, that is, one numbered less than 1024. |
| <code>secure_locks</code> | Requires that clients requesting lock operations be properly authenticated before activating the lock (synonym for <code>auth_nlm</code>). |
| <code>subtree_check</code> | If only part of a file system, such as a subdirectory, is exported, subtree checking makes sure that file requests apply to files in the exported portion of the file system. |
| <code>sync</code> | Forces the server to perform a disk write before notifying the client that the request is complete. |
| <code>wdelay</code> | Instructs the server to delay a disk write if it believes another related disk write may be requested soon or if one is in progress, improving overall performance. |

b Write the use of `/etc/sysconfig/init` file.

The `/etc/sysconfig/init` file controls how the system will appear and function during the boot process.

The following values may be used:

BOOTUP=*value*, where *value* is one of the following:

BOOTUP=color means the standard color boot display, where the success or failure of devices and services starting up is shown in different colors.

BOOTUP=verbose means an old-style display, which provides more information than purely a message of success or failure. Anything else means a new display, but without ANSI formatting.

RES_COL=*value*, where *value* is the number of the column of the screen to start status labels. It defaults to 60.

MOVE_TO_COL=*value*, where *value* moves the cursor to the value in the `RES_COL` line. It defaults to ANSI sequences output by `echo -e`.

SETCOLOR_SUCCESS=*value*, where *value* sets the color to a color indicating success. It defaults to ANSI sequences output by `echo -e`, setting the color to green.

SETCOLOR_FAILURE=*value*, where *value* sets the color to one indicating failure. It defaults to ANSI sequences output by `echo -e`, setting the color to red.

SETCOLOR_WARNING=*value*, where *value* sets the color to one indicating warning. It defaults to ANSI sequences output by `echo -e`, setting the color to yellow.

SETCOLOR_NORMAL=*value*, where *value* sets the color to “normal.” It defaults to ANSI sequences output by `echo -e`.

LOGLEVEL=*value*, where *value* sets the initial console logging level for the kernel. The default is 7; 8 means everything (including debugging); 1 means nothing except kernel panics. `syslogd` will override this once it starts.

PROMPT=*value*, where *value* is one of the following Boolean values: `yes` — Enables the key check for interactive mode. `no` — Disables the key check for interactive mode.

c What is NFS? [1 mark]

What are the features of NFS? [2 mark]

Enumerate the additional features of NFS4. [2 mark]

NFS follows standard client/server architectural principles. The server component of NFS consists of the physical disks that contain the file systems you want to share and several daemons that make these shared file systems visible to and available for use by client systems

on the network. When an NFS server is sharing a file system in this manner, it is said to be *exporting a file system*. Similarly, the shared file system is referred to as an *NFS export*.

Features of NFSv4

- NFSv4 incorporates RPCSEC-GSS (the SecureRPC protocol using the Generic Security Service API) security, which makes it possible to encrypt the data stream transmitted between NFS clients and servers.
- Another security feature added to NFSv4 is support for access control lists, or ACLs.
- In terms of performance enhancements, NFSv4 makes fuller use of clientside caching, which reduces the frequency with which clients must communicate with an NFS server. By decreasing the number of server round trips, overall performance increases.
- In addition, NFSv4 was specifically designed (or enhanced) to provide reasonable performance over the Internet, even on slow, low-bandwidth connections or in high latency situations
- Complementing the new version's greater Internet-friendliness, NFSv4 also supports Unicode (UTF-8) filenames, making cross-platform and intercharacter set file sharing more seamless and more international.

d Explain the concept of supernetting with suitable example.

Under supernetting, the class subnet masks are extended so that a network address and subnet mask could, for example, specify multiple Class C subnets with one address. For example, if you needed about a thousand addresses, you could supernet four Class C networks together:

192.60.128.0 (11000000.00111100.10000000.00000000) Class C subnet address

192.60.129.0 (11000000.00111100.10000001.00000000) Class C subnet address

192.60.130.0 (11000000.00111100.10000010.00000000) Class C subnet address

192.60.131.0 (11000000.00111100.10000011.00000000) Class C subnet address

192.60.128.0 (11000000.00111100.10000000.00000000) Supernetted Subnet address

255.255.252.0 (11111111.11111111.11111100.00000000) Subnet Mask

192.60.131.255 (11000000.00111100.10000011.11111111) Broadcast address

In this example, the subnet 192.60.128.0 includes all the addresses from 192.60.128.0 to 192.60.131.255. As you can see in the binary representation of the subnet mask, the network portion of the address is 22 bits long, and the host portion is 10 bits long.

Under CIDR, the subnet mask notation is reduced to simplified shorthand. Instead of spelling out the bits of the subnet mask, the number of 1 bits that start the mask are simply listed. In the example, instead of writing the address and subnet mask as

192.60.128.0, Subnet Mask 255.255.252.0

the network address is written simply as

192.60.128.0/22

This address indicates the starting address of the network, and number of 1 bits (22) in the network portion of the address. If you look at the subnet mask in binary, you can easily see how this notation works.

(11111111.11111111.11111100.00000000)

3 Attempt any two of the following:

10

a Explain the process of connecting to samba client from Linux.

The connection can be made from the command line using two methods. The first uses a utility called smbclient, and the command syntax is smbclient //computer name/sharename, as shown in the following example. Be sure to replace the computer name in the example with the name of your computer.

```
[root@terry terry]# smbclient //terrycollings/c
```

```
added interface ip=192.168.9.93 bcast=192.168.9.255 nmask=255.255.255.0
```

```
Got a positive name query response from 192.168.9.102 (192.168.9.102)
```

```
Password:
```

```
Domain=[Tardis] OS=[Windows 5.0] Server=[Windows 2000 LAN Manager]
```

```
smb: \>
```

The preceding example shows me logging in to my Windows PC from my Red Hat system. I was prompted for a password to log in and was given some information about the Windows system and a command prompt. You can type help at the command prompt to get a list of possible commands. The commands at the smb prompt are very similar to command-line FTP commands.

To exit the connection, type **exit**.

Another way to make the files on the Samba client accessible on your Red Hat system is to mount the client file system on your file system. You can do this using the smbmount command. The syntax for this command is smbmount //computer name/directory /mysystem/mount/point, as shown in the following example:

```
[root@terry terry]# smbmount //terrycollings/c /mnt/windows
```

Password:

Next, you can change to the directory on the system where you mounted the Windows system by issuing the following command:

```
[root@terry terry]# cd /mnt/windows
```

Then you can get a directory listing by using the ls command.

You can put the mount command into a local startup script so that the directories are mounted at system boot, if you desire. Use the command as shown earlier and add an option to look for a file that contains the login username and password.

```
smbmount //terrycollings/c /mnt/windows -o credentials=/home/terry/.sambacred
```

You need to create a file as shown in the following code. I created a hidden file called .sambacred and in the file I placed these two lines:

```
Username = terry
```

```
password = (password)
```

- b List any five NTP utility programs and write the purpose of each. [1 mark each]

| PROGRAM | DESCRIPTION |
|------------|--|
| ntpd | Sets the system date and time via NTP |
| ntpd | Controls the NTP daemon, ntpd |
| ntp-keygen | Generates public and private keys for use with NTP |
| ntpq | Queries the NTP daemon |
| ntpsim | Provides NTP simulation for development and testing |
| ntpdate | Displays the time variables maintained by the Linux kernel |
| ntptrace | Traces a chain of NTP servers back to the primary source |
| tickadj | Sets certain time variables maintained by the Linux kernel |

- c Write the purpose of home section of samba configuration file.

[homes], is used to enable the server to give users quick access to their home directories.

comment = Home Directories — A comment line.

browseable = yes — Means that the directory will appear in the Windows file browser.

writable = yes — Means that users can write to their directories.

create mode = 0664—Sets the default file permissions for files created in the directory.

directory mode = 0775 — Sets the default permissions for created directories.

max connections = 1 — The maximum number of simultaneous connections allowed. Setting this number to 1 prevents a user from logging in to the server from more than one location. Setting this number to 2 allows a user to log in from two locations and so on. Setting this number

to 0 allows an unlimited number of connections.

- d Write the steps to configure squid.

The initialization script that controls Squid is

/etc/rc.d/init.d/squid, which reads default values from /etc

/sysconfig/squid.

Table lists the configuration settings with which you concern yourself in the short term.

As Table shows, cache_effective_user and cache_effective_group identify the user ID (UID) and group ID (GID), respectively, under which Squid runs. The default values, squid, are the

defaults configured in the Squid package shipped in Fedora Core and RHEL. You needn't change them.

`httpd_accel_with_proxy`, which defaults to off, controls whether Squid runs as a cache (or accelerator) and proxy or just as proxy. When set to off, Squid functions only as a proxy. If set to on, Squid works as both a cache *and* a proxy. If you are using Squid's caching functionality, you'll need to set `httpd_accel_port` to 80 and use `httpd_accel_host` to define the name of the host running Squid. As it happens, the default port number of `httpd_accel_port` is 80, so you shouldn't need to change this. `httpd_accel_host` lacks a default value, so you'll need to change this value. If you want a transparent proxy server, set `httpd_accel_uses_host_header` to on. The default value, off, means that clients have to configure their Web clients to use a proxy server, which can be quite inconvenient for users and a pain for administrators to manage, especially across a LAN that is geographically dispersed or if some users are, shall we say, technically challenged.

The final value to configure is `httpd_access`, which controls who can access the Squid server and, therefore, who can surf the Web through the proxy. The default configuration is deny all, which prevents *any* user from accessing the proxy. As you can imagine, this policy is draconian. For experimentation purposes, set it to allow all, which permits all users to access the server. Yes, this is just as extreme as deny all, but the idea is to enable people to surf the Web. Make sure to learn how to use Squid's access control list (ACL) features to fine-tune the access policy.

The following listing shows the changes to make to `/etc/squid/squid.conf`:

```
cache_effective_user squid
cache_effective_group squid
httpd_accel_host squid.example.com
httpd_accel_with_proxy on
httpd_accel_port 80
httpd_accel_uses_host_header on
httpd_access allow all
```

Replace `squid.example.com` with the name of the system on which you are running Squid. To save needing to do so later, initialize Squid's cache using the command `squid -z`:

```
# squid -z
```

```
2005/03/02 22:54:59| Creating Swap Directories
```

4 Attempt *any two* of the following:

10

a What is SSH? [1 mark]

Explain with suitable example. [4 mark]

Secure Shell, also known as SSH, is a secure Telnet replacement that encrypts all traffic, including passwords, using a public/private encryption key exchange protocol. It provides the same functionality of Telnet, plus other useful functions, such as traffic tunneling.

This is what it looks like to SSH into a machine for the first time:

```
[vnavrat@buffy vnavrat$ ssh vnavrat@woolf.xena.edu
The authenticity of host 'woolf.xena.edu (123.456.789.65)'
can't be established.
```

```
RSA key fingerprint is
```

```
b2:60:c8:31:b7:6b:e3:58:3d:53:b9:af:bc:75:31:63.
```

```
Are you sure you want to continue connecting (yes/no)? yes
```

```
Warning: Permanently added 'woolf.xena.edu,123.456.789.65'
(RSA) to the list of known hosts.
```

```
vnavrat@woolf.xena.edu's password:
```

```
Welcome to woolf
```

```
Unauthorized usage prohibited. Please check your quotas.
```

```
vnavrat:~>
```

Following is an example of how to tunnel your mail through SSH to keep your password and mail encrypted and secure during transit. In this example, you use POP3 to retrieve your mail

from the remote machine buffy.xena.edu. Normally you would tell your POP3 software to connect from your localhost to port 110 (the POP port) of buffy.xena.edu.

But in this example the first step is to configure your POP mailer to connect to port 16510 of your own machine, and put in the password for your account on buffy.xena.edu. The second step is to set up the SSH tunnel, which encrypts and forwards the traffic over the network to terry.muhlenberg.edu's POP port.

To set up the SSH tunnel, type the following at the command line:

```
ssh -N -L 16510:127.0.0.1:110 terry@terry.muhlenberg.edu
```

b List [1 mark]

and explain different types of domain name servers. [4 mark]

Master — The master contains all the information about the domain and supplies this information when requested. A master server is listed as an authoritative server when it contains the information you are seeking and it can provide that information.

Slave— The slave is intended as a backup in case the master server goes down or is not available. This server contains the same information as the master and provides it when requested if the master server cannot be contacted.

Caching— A caching server does not provide information to outside sources; it is used to provide domain information to other servers and work stations on the local network. The caching server remembers the domains that have been accessed. Use of a caching server speeds up

searches since the domain information is already stored in memory, and the server knows exactly where to go rather than having to send out a request for domain information.

c Explain the use of reverse zone files with suitable example.

This file provides information to map IP addresses to names.

You can also find a name from an IP number, and this is called reverse address resolution. All you need to do is enter the IP address, and the server returns the domain name. Reverse address resolution requires the use of a reverse zone file.

Following is a sample reverse zone file.

```
@      IN      SOA      localhost. root.localhost. (
        1997022700 ; Serial
        28800      ; Refresh
        14400      ; Retry
        3600000    ; Expire
        86400     ) ; Minimum
1      IN      NS      localhost.
1      IN      PTR     localhost.
```

d List various less secure services. [1 mark]

Write the purpose of any four. [1 mark each]

- telnet
- ftp
- rsync
- rlogin
- rsh
- finger
- talk and ntalk

5 Attempt any two of the following:

a List various protocols used to fetch the mail from mail server. [1 mark]

Explain working of each in short. [4 mark]

POP3 Post Office Protocol version 3

IMAP4 Internet Message Access Protocol version 4

POP3 runs on a server that is connected to a network and that continuously sends and receives mail. The POP3 server stores any messages it receives until the message recipients request them POP3 uses the MTA's storage to hold messages until they are requested. When users want to check their email, they connect to the POP3 server and retrieve messages that were

stored by the server. After retrieval, the messages are stored locally (that is, locally to the MUA) and you can use the MUA on your PC to read them at your leisure. Of course, your MUA has to understand the POP3 to be able to communicate with the POP3 server, but most MUAs speak fluent POP3 these days so this is rarely a problem. The messages you retrieve to your PC are then typically removed from the server.

What happens if users don't want to remove their email from the server and instead want to access their email from any given PC?

The Internet Message Access Protocol version 4 (IMAP4) provides much more sophisticated email-handling functionality than SMTP or POP3 do. IMAP4 has more features. IMAP4 enables you to store email on a networked mail server, just as POP3 does. The difference is that POP3 requires you to download your email before your MUA reads it, whereas IMAP4 enables your email to reside permanently on a remote server, from which you can access your mail.

And

you can do so from your office, your home, your PDA, your cell phone, or anywhere else. Your MUA must understand IMAP4 to retrieve messages from an IMAP4 server.

b Write several key components that are essential for email to work properly.

Several key components are essential for email to work properly

Programs:

A mail user agent for users to be able to read and write email

A mail transfer agent to deliver email messages between computers across a network

A mail delivery agent to deliver messages to users' mailbox files

A mail-notification program to tell users that they have new mail (optional)

The SMTP protocols for packaging email and transferring email messages between MTAs

c Write the purpose of the following parameters of vsftpd.conf file [1 mark each]

a. `anonymous_enable` : `anonymous_enable=YES`, allows anonymous FTP access. You can set this to `NO` if you do not want to enable anonymous FTP.

b. `write_enable` : `write_enable=YES` enables all variations of the FTP commands that allow FTP users to modify the file system, such as `STOR` and `DELE`. By setting `write_enable` to `NO` you can disable the write commands.

c. `chown_username` : specifies the name of the user to set ownership of uploaded files

d. `ftpd_banner` : allows you to display a site-specific banner message when users connect to the server.

e. `dirmessage_enable` : If this directive is set to `YES`, the first time a user enters a new directory, vsftpd displays the contents of a file named `.message`, if it exists.

d Explain how to disable anonymous FTP.

The easiest way is to remove the ftp user from `/etc`

`/passwd` and `/etc/group`:

```
# cp -p /etc/passwd /etc/passwd.ftp
```

```
# cp -p /etc/group /etc/group.ftp
```

```
# userdel -r ftp
```

```
userdel: /var/ftp not owned by ftp, not removing
```

```
# find / -user 50 | xargs rm -r
```

Ordinarily, `userdel`'s `-r` option removes files in ftp's home directory (`/var/ftp`), but it doesn't work in this case because the ftp user doesn't own `/var/ftp`, root does. `userdel` also removes the ftp user from `/etc/group`, so you needn't execute the `groupdel` command. The `find` command locates all the files owned by the ftp user and deletes them. You have to use the numeric UID (50) instead of the username (ftp) because the username no longer exists. You might not want to execute the command if you have populated the FTP server with files that you can't easily replace.

The problem with this method is that if you later decide to permit anonymous FTP, you have to recreate the ftp user and group because, as configured, vsftpd doesn't allow *any* FTP login if the user ftp is not present in the password file. That's why we made backup copies of `/etc/passwd` and `/etc/group` before executing `userdel`.

A more flexible approach is to add ftp to /etc/vsftpd/user_list and set userlist_deny=YES and anonymous_enable=NO in /etc/vsftpd/vsftpd.conf. It is *not* sufficient to comment out anonymous_enable=YES, because that will default to permitting anonymous FTP. This approach disables anonymous FTP while permitting regular FTP. However, if you use this method, remove any other users from /etc/vsftpd/user_list that you do want to be able to log in via FTP.

6 Attempt any two of the following:

a Write the purpose of any five Global Configuration Directives of httpd.conf. [1 mark each]

| DIRECTIVE | DESCRIPTION |
|-----------------------------|---|
| Include conf.d/*.conf | Includes the contents of the files in conf.d/ whose names end in .conf |
| KeepAlive Off | If set to On, maintains an open connection to a remote client in the absence of direct contact for the time specified by KeepAliveTimeout |
| KeepAliveTimeout 15 | Sets the number of seconds permitted to elapse between direct requests from the same client on the same connection before the server will close the connection (applies if KeepAlive is On) |
| Listen [ipaddress:]80 | Determines the combination of IP address and port on which Apache listens for connections; multiple Listen directives may be used |
| LoadModule modname filename | Links the module or library filename into the server and adds it to the list of active modules using the name modname |
| MaxClients 256 | Sets the maximum number of simultaneous connections supported |
| MaxKeepAliveRequests 100 | Sets the number of requests permitted per connection |
| MaxRequestsPerChild 4000 | Sets the maximum number of requests each child server fills before terminating |
| MaxSpareServers 20 | Defines the maximum number of spare (idle) child servers the master server spawns |
| MinSpareServers 5 | Defines the minimum number of spare (idle) child servers permitted |
| PidFile run/httpd.pid | Defines the file containing the PID of the master server process, relative to ServerRoot |
| ServerLimit 256 | Specifies the upper limit on the number of server processes or threads running simultaneously |
| ServerRoot /etc/httpd | Defines the top-level directory for Apache's configuration files and log files |
| ServerTokens OS | Defines the contents of the server's HTTP response header |
| StartServers 8 | Defines the number of child servers created when Apache starts |
| Timeout 120 | Defines the maximum time in seconds Apache waits for packet send and receive operations to complete |

b Write an SSI page that will display the long listing of the directory /home/tyit.

```
<html>
<head>
<title>SSI Test Page</title>
<link rel="stylesheet" type="text/css" href="rhlnsa3.css">
</head>
<body>
<h1>SSI Test Page</h1>
<div id="content">
<pre>
<!--#exec cmd="ls -l /home/tyit" -->
</pre>
</div> <!-- content -->
<!--#include virtual="footer.html" -->
</body>
```

</html>

<Directory “/var/www/html/tests”>

Options Indexes FollowSymLinks Includes

AllowOverride None

Order allow,deny

Allow from all

</Directory>

c What is root account?

What is sudo?

What are its features?

How does a sudo session work?

The root account has unlimited power on any Linux or UNIX system, and, in this respect, Red Hat Linux is no exception.

Sudo enables you to give specific users or groups of users the ability to run some (or all) commands

requiring root privileges. Sudo also logs all commands executed, which allows you to maintain an audit trail of the commands executed, by whom they were executed, when they were executed, and so on.

Sudo’s features include:

- Enabling the ability to restrict the commands a given user may run on a per-host basis.

- Maintaining a clear audit trail of who did what. The audit trail can use the system logger or Sudo’s own log file. In fact, you can use Sudo in lieu of a root shell to take advantage of this logging.

- Limiting root-equivalent activity to a short period of time using timestamp based “tickets,” thus avoiding the potential of leaving an active root shell open in environments where others can physically get to your keyboard.

- Allowing a single configuration file, /etc/sudoers, to be used on multiple machines, permitting both centralized Sudo administration and the flexibility to define a user’s privileges on a per host basis.

Sudo session proceeds

as follows:

1. An authorized user prefixes the root command she wants to execute with sudo followed by a space, for example:

\$ sudo shutdown -h +5 “System shutting down for disk replacement”

2. Sudo prompts the user for her personal password (*not* the root password) and then checks the configuration file (/etc/sudoers) to make sure she has permission to run the given command on a given machine. The password prompt can be overridden by specifying the NOPASSWD flag in /etc/sudoers, but this poses as security risk, so we don’t cover the NOPASSWD flag in this section.

3. If the user is permitted to use that command, Sudo runs the command as root (or another user if specified), logs the details, and timestamps the Sudo session ticket.

4. If the user is *not* permitted to use that command, Sudo logs the attempt and exits. Sudo also logs problems and other invalid sudo uses.

5. After executing the first command, the user can use multiple sudo commands without being prompted for her password again. The session ticket expires five minutes (the default expiration period) after the

last sudo command is issued, after which the user is again prompted for a password.

d Explain the use of rpmquery commands.

The general form of an RPM query is:

rpmquery [query_opts]

rpmquery (or, if you prefer the old style, rpm -q or rpm --query) specifies a query operation and query_opts specifies what to query, the type of query, how the query should run, or the format

of its output. You can use the command `rpmquery` in place of `rpm -q` or `rpm --query`. Most commonly,

queries use the following general syntax:

```
rpmquery [query_opts] package [...]
```

`package` names the RPM to query. Query multiple RPMs using a spaceseparated list of package names. Query mode's power comes at the cost of a long list of options for the `query_opts` argument. The options fall into two broad categories. One group, referred to as *package selection options*, controls which package or packages to query, and the other, known as *output selection options*, defines what information to display.

Table lists many but not all of the options available in query mode.

The Type column uses S to mark a package selection option and I to mark an information selection option.

| | | |
|--|---|--|
| <code>-a</code> | S | Queries all installed RPMs. Does not require a package specification. |
| <code>-c</code> | I | Lists only the configuration files stored in the queried RPM(s). |
| <code>--changelog</code> | I | Displays change information about the queried RPM(s). |
| <code>-d</code> | I | Lists only the documentation files stored in the RPM. |
| <code>--dump</code> | I | For each file stored in the queried RPM(s), displays its path, size, modification time, MD5 checksum, permissions, owner, group, and whether it is a configuration file, documentation file, a device, or a symlink (must be used with <code>-l</code> , <code>-c</code> , or <code>-d</code>). |
| <code>-f file</code> | S | Queries the RPM that owns <code>file</code> . Does not require a package specification. |
| <code>-g group</code> | S | Lists the packages in the RPM group named <code>group</code> . Does not require a package specification. |
| <code>-i</code> | I | Displays complete information about the queried RPM(s). |
| <code>-l</code> | I | Lists all of the files stored in the RPM. |
| <code>--last</code> | I | Displays the installation date and time of each RPM queried, starting with the most recently installed RPM. |
| <code>-p package [...]</code> | S | Queries the uninstalled RPM named <code>package</code> . |
| <code>--provides</code> | I | Lists all of the capabilities the queried RPM(s) provides. |
| <code>--qf 'format_str'</code> | I | Creates a customized output format for displayed information, using <code>format_str</code> as the model. |
| <code>--querytags</code> | I | Prints all known tags for use with the <code>--qf</code> option. Does not require a package specification. |
| <code>--requires</code> | I | Lists all RPMs on which the package depends. |
| <code>-s</code> | I | For each file in the original RPM, displays its state, which is one of normal, not installed, or replaced. |
| <code>--whatprovides capability</code> | S | Queries all RPMs that provide <code>capability</code> . |
| <code>--whatrequires capability</code> | S | Queries all RPMs that need <code>capability</code> in order to function properly. |

7 Attempt any three of the following:

a Explain the different runlevels in Linux.

0 — Halt

1 — Single-user mode

2 — Not used (user-definable)

3 — Full multiuser mode (without a graphical user interface, GUI)

4 — Not used (user-definable)

5 — Full multiuser mode (with a GUI)

6 — Reboot

b What is DHCP? Explain its configuration file.

Using Dynamic Host Configuration Protocol (DHCP), you can have an IP address and the other information automatically assigned to the hosts connected to your network.

The configuration file for DHCP is /etc/dhcpd.conf

```
 #(The amount of time in seconds that the host can keep the IP address.)
```

```
 default-lease-time 36000;
```

```
 #(The maximum time the host can keep the IP address.)
```

```
 #domain name
```

```
 max-lease-time 100000;
```

```
 # (The domain of the DHCP server.)
```

```
 #nameserver
```

```
 option domain-name "tactechology.com";
```

```
 option domain-name-servers 192.168.1.1;
```

```
 #gateway/routers, can pass more than one:
```

```
 option routers 1.2.3.4,1.2.3.5;
```

```
 option routers 192.168.1.1; (IP address of routers.)
```

```
 #netmask (The subnet mask of the network.)
```

```
 option subnet-mask 255.255.255.0;
```

```
 #broadcast address (The broadcast address of the network.)
```

```
 option broadcast-address 192.168.1.255;
```

```
 #specify the subnet number gets assigned in
```

```
 subnet 192.168.1.0 netmask 255.255.255.0
```

```
 #define which addresses can be used/assigned
```

```
 range 192.168.1.1 192.168.1.126;
```

c Explain smbclient and smbmount commands of samba.

The syntax for this command is

```
smbmount //computer name/directory /mysystem/mount/point
```

example:

```
[root@terry terry]# smbmount //terrycollings/c /mnt/windows
```

Password:

Command syntax is smbclient //computer name/sharename

Example:

```
[root@terry terry]# smbclient //terrycollings/c
```

```
added interface ip=192.168.9.93 bcast=192.168.9.255 nmask=255.255.255.0
```

```
Got a positive name query response from 192.168.9.102 (192.168.9.102)
```

Password:

```
Domain=[Tardis] OS=[Windows 5.0] Server=[Windows 2000 LAN Manager]
```

```
smb: \>
```

d Explain the concept of domain with suitable example.

Ip address is expressed as a group of numbers referred to as a dotted quad group. These groups of numbers present no problem to the computers in the network, but it is difficult for humans to remember many groups of numbers. So, you need to be able to enter names and then have these names converted into numbers. Each time you type a Web site's address into your browser, the Domain Name System (DNS) goes to work. You enter names that are easy for you to remember, and the names are resolved into numbers that computers find easy to understand. Enabling efficient human/machine interaction is the function of name address resolution. In this chapter you learn how to install and configure the Domain Name System, which provides this name address resolution.

First, take a look at domain names and their organization using the domain name tactechology.com. The first part of this domain name, tactechology, is the name of the company, institution, or organization. The next part, after the period (dot in today's vernacular) is called the top-level domain (TLD). In addition to the com top-level domain, you will find many others.

| TOP-LEVEL DOMAIN | MEANING |
|------------------|--|
| com | A TLD typically used to register a business |
| edu | A TLD assigned to an educational institution |
| gov | A TLD given to a U.S. government agency |
| mil | A TLD used by a branch of the U.S. military |
| net | A TLD used by a network affiliated organization |
| org | A TLD used by a noncommercial organization |
| int | An TLD assigned to an international organization |
| us | The U.S. domain, with each listing as a lower level |
| biz | Another TLD used for businesses |
| info | A TLD that can be used for any type of organization |
| name | A TLD used to register sites for individuals |
| museum | A TLD used to register a museum |
| coop | A TLD assigned to a cooperative organization |
| aero | A TLD used by the air transport industry |
| pro | A TLD not yet active but assigned to professions |
| travel | A TLD that should be available in late 2005 used by travel related companies |

Large domains may be further broken down into subdomains. For example, the tactechology site is `www.tactechology.com`. Perhaps the accounting department runs their own Web server. To find their Web server, tactechology contains the subdomain `acct.tactechology.com`. An individual computer in the accounting department also has a hostname, for example, `payables`. The complete name for this computer is then `payables.acct.tactechology.com`, and you can find its IP address by using the DNS to look it up.

e Write the purpose of `ftpusers` and `user_list` files of `ftp`.

f. Explain `useradd` command with suitable example.

The `useradd` command creates new user accounts and, when invoked with the `-D` option, modifies the default values applied to new accounts. As a result, it can be invoked in two ways. The syntax of the first form is:

```
useradd [-c comment] [-d dir] [-e date]
[-f time] [-g initial] [-G group[,...]]
[-m [-k dir] | -M]
[-p passwd] [-s shell] [-u uid [-o]]
[-n] [-r] username
```

The first form creates a new user account named `username`. Optional values not specified using options are assigned default values drawn from `/etc/login.defs` and `/etc/default/useradd`. Table lists the options `useradd` accepts.

| OPTION | DESCRIPTION |
|---------------------------------------|---|
| <code>-c <i>comment</i></code> | Uses <code>comment</code> for the name field |
| <code>-d <i>dir</i></code> | Names the new user's home directory <code>dir</code> |
| <code>-e <i>date</i></code> | Sets the account's expiration date to <code>date</code> |
| <code>-f <i>time</i></code> | Disables the account <code>time</code> days after the password expires |
| <code>-g <i>group</i></code> | Sets the user's primary group membership, or login group, to <code>group</code> |
| <code>-G [<i>group</i>[, ...]]</code> | Makes the user a member of each supplemental group <code>group</code> |

| | |
|------------------------|---|
| <code>-m</code> | Creates the home directory if it does not exist and copies the files and directory structure in <code>/etc/skel</code> to the new directory |
| <code>-k dir</code> | Copies the files and directory structure in <code>dir</code> , not <code>/etc/skel</code> , to the new home directory; must be specified with <code>-m</code> |
| <code>-M</code> | Disables creation of the home directory; cannot specify <code>-m</code> and <code>-M</code> |
| <code>-p passwd</code> | Sets the account password to the encrypted password <code>passwd</code> |
| <code>-s shell</code> | Sets the user's default shell to <code>shell</code> |
| <code>-u uid</code> | Sets the user's UID (User ID) to <code>uid</code> , which must be a unique number |
| <code>-o</code> | Allows the UID specified with <code>-u uid</code> not to be unique; must be specified with <code>-u</code> |
| <code>-n</code> | Disables use of Red Hat's user private groups |
| <code>-r</code> | Creates a system account (an account with a UID less than 100) but does not create a home directory |
| <code>username</code> | Sets the login name to <code>username</code> |

The second way to invoke `useradd` uses the `-D` option. Invoked with only `-D`, `useradd` displays its current default settings. Using `-D` with any of the options listed in Table modifies the default value for the corresponding field. Here is the syntax for the second form:

```
useradd -D [-g group] [-b home_dir]
[-f inactive_time] [-e expire_date]
[-s shell]
```

`useradd`'s default values are stored in `/etc/default/useradd`.

| OPTION | DESCRIPTION |
|-----------------------|---|
| <code>-g group</code> | Sets the default group to <code>group</code> |
| <code>-b dir</code> | Sets the default home directory to <code>dir</code> |
| <code>-f time</code> | Sets the default account disable time to <code>time days</code> |
| <code>-e date</code> | Sets the default account expiration date to <code>date</code> |
| <code>-s shell</code> | Sets the default login shell to <code>shell</code> |

