DUMPSARENA

# RHCE (Redhat Certified Engineer)

## RedHat RH302

**Version Demo**

**Total Demo Questions: 39**

**Total Premium Questions: 397**

**Buy Premium PDF**

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

# Topic Break Down

| Topic | No. of Questions |
|---|---|
| Topic 1, System Configuration and Management | 149 |
| Topic 2, Network Services | 169 |
| Topic 3, Security | 72 |
| Topic 4, Mix Questions | 7 |
| Total | 397 |

QUESTION NO: 1 - (SIMULATION)

SIMULATION

Make a swap partition having 100MB. Make Automatically Usable at System Boot Time.

**ANSWER: See the explanation for the answer**

**Explanation:**

Create a new **100 MiB swap partition** on the available unused disk space, initialize it as swap, enable it now, and add it to `/etc/fstab` so it is activated automatically at boot.

```
# Identify the target disk and free space
lsblk
fdisk -l

# Replace /dev/sdX with the disk that has unused space
fdisk /dev/sdX
```

In `fdisk`, create a new partition of exactly 100 MiB and set its type to Linux swap:

```
n                 # new partition
# accept the appropriate/default partition number and first sector
+100M             # last sector / size

t                 # change partition type (if it was not selected automatically)
# select the new partition number
82                # Linux swap type on DOS/MBR partition tables
w                 # write changes and exit
```

For GPT disks, select the partition type named `Linux swap` (fdisk type code is commonly `8200`) rather than `82`.

```
# Inform the kernel of the changed partition table (replace sdXn appropriately)
partprobe /dev/sdX

# Create the swap area and enable it immediately
mkswap /dev/sdXn
swapon /dev/sdXn

# Obtain its UUID
blkid /dev/sdXn
```

Add a persistent mount entry to `/etc/fstab`, using the UUID returned by `blkid`:

```
UUID=<uuid-of-/dev/sdXn>  none  swap  defaults  0 0
```

Validate the configuration:

```
swapon --show
free -m
swapoff /dev/sdXn && swapon -a
swapon --show
```

The partition is now active and the `/etc/fstab` entry ensures it will be enabled automatically during system boot.

**QUESTION NO: 2 - (SIMULATION)**

SIMULATION

Create the directory /archive and group owner should be the sysuser group.

**ANSWER: See the explanation for the answer**

**Explanation:**

Create the directory, then set its group ownership to `sysuser`:

```
mkdir /archive
chgrp sysuser /archive
```

Equivalently, in one ownership command:

```
mkdir /archive
chown :sysuser /archive
```

Verify the result:

```
ls -ld /archive
```

The group column in the output should show `sysuser`, for example:

```
drwxr-xr-x. 2 root sysuser ... /archive
```

## QUESTION NO: 3 - (SIMULATION)

**/data directory on linux server should make available on windows system that eric user should able to access on read only mode within example.com domain.**

## ANSWER: See the explanation for the answer

**Explanation:**

Configure Samba to export `/data` as a read-only share and permit only the Samba user `eric`.

```
dnf install -y samba samba-client
mkdir -p /data
setfacl -m u:eric:rx /data
semanage fcontext -a -t samba_share_t '/data(/.*)?'
restorecon -RFv /data
setsebool -P samba_export_all_ro on
smbpasswd -a eric
```

Add (or amend) the following settings in `/etc/samba/smb.conf`:

```
[global]
        workgroup = EXAMPLE
        security = user

[data]
        path = /data
        browseable = yes
        read only = yes
        valid users = eric
```

`workgroup = EXAMPLE` corresponds to the Windows/Samba domain name for `example.com`. If the server is already joined to an AD domain, retain its existing global AD/domain configuration and add only the `[data]` share; use the domain identity in `valid users` if required, for example `valid users = EXAMPLE\eric`.

Validate the configuration, enable Samba, and permit SMB through the firewall:

```
testparm
systemctl enable --now smb nmb
firewall-cmd --permanent --add-service=samba
firewall-cmd --reload
```

From Windows, Eric connects to `\\<server-name>\data` (or `\\<server-IP>\data`) and authenticates as `eric`. The Samba `read only = yes` setting prevents writes even if filesystem permissions would otherwise allow them.

## QUESTION NO: 4 - (SIMULATION)

**Install the Redhat Linux RHEL 5 through NFS. Where your Server is server1.example.com having IP 192.168.0.254 and shared /var/ftp/pub. The size of the partitions are listed below:**

**/?1048**

**/home?1028**

**/boot? 512**

**/var? 1028**

**/usr?2048**

**Swap->1.5 of RAM Size**

**/data?configure the RAID Level 0 of remaining all free space.**

**After completing the installation through NFS solve the following questions. There are two networks 192.168.0.0/24 and 192.168.1.0/24. As well as there are two domains example.com on 192.168.0.0/24 network and cracker.org on 192.168.1.0/24 network. Your system is based on example.com domain.**

## ANSWER: See the explanation for the answer

**Explanation:**

Boot the RHEL 5 installation media (or PXE boot media) and select **NFS installation**. Configure the installer network interface with an unused address on `192.168.0.0/24` (and the appropriate netmask/gateway for that lab). The installation source is:

```
NFS server: 192.168.0.254
NFS directory: /var/ftp/pub
```

Equivalently, the NFS installation location is `nfs:192.168.0.254:/var/ftp/pub`.

In the manual partitioning screen, create standard Linux partitions/filesystems with these sizes (the values are normally entered as MB):

`/boot` — `512 MB`, `ext3`. Do not place `/boot` on the RAID 0 device.
`/` — `1048 MB`, `ext3`.
`/home` — `1028 MB`, `ext3`.
`/var` — `1028 MB`, `ext3`.
`/usr` — `2048 MB`, `ext3`.
`swap` — size equal to `1.5 × the installed RAM`. For example, with 1024 MB RAM, create a 1536 MB swap partition.

Use **all disk space left after the above allocations** for RAID members. Create at least two partitions of type `fd` (Linux RAID autodetect), normally one on each available disk, and use all remaining space. Then create a software RAID device with:

```
RAID device: /dev/md0
RAID level: RAID 0 (stripe)
Mount point: /data
Filesystem: ext3
```

Assign every remaining RAID member partition to `/dev/md0`; therefore `/data` consumes all remaining free space. RAID 0 provides capacity/performance but **no fault tolerance**. If the machine has only one disk, RAID 0 can only be created by splitting that disk into multiple RAID-member partitions, which is not useful for redundancy; normally this task assumes at least two disks.

Continue the installation, install the boot loader to the normal boot disk, and set the installed system hostname in the `example.com` domain (for example, `client.example.com`, using the hostname assigned by the lab). The stated `cracker.org` network is not used for this NFS installation; it is likely relevant to subsequent networking/DNS tasks.

## QUESTION NO: 5 - (SIMULATION)

SIMULATION

If any mail coming from outside of the local LAN block all mails.

## ANSWER: See the explanation for the answer

**Explanation:**

Configure Sendmail's access database to allow only the local LAN and reject every other SMTP client connection. For example, if the local LAN is `192.168.1.0/24`, add the following to `/etc/mail/access` (use the actual local LAN prefix):

```
Connect:127.0.0.1          RELAY
Connect:192.168.1          RELAY
Connect:ALL                REJECT
```

`Connect:192.168.1` matches clients in `192.168.1.0/24`. The final `Connect:ALL REJECT` is the catch-all rule and denies SMTP connections from any source not matched by the preceding local rules.

Rebuild the Sendmail access map and restart Sendmail:

```
makemap hash /etc/mail/access.db < /etc/mail/access
systemctl restart sendmail
```

On older SysV-based systems, use:

```
service sendmail restart
chkconfig sendmail on
```

Do not use a `From:` rule for this requirement: `Connect:` rejects the remote SMTP host itself, ensuring all mail from outside the LAN is blocked.

## QUESTION NO: 6 - (SIMULATION)

**Install the Cron Schedule for david user to display “Hello” on daily 5:30.**

## ANSWER: See the explanation for the answer

**Explanation:**

Install a per-user crontab entry for `david` that runs every day at 05:30:

```
crontab -u david -e
```

Add this line:

```
30 5 * * * /usr/bin/echo "Hello"
```

Alternatively, non-interactively:

```
echo '30 5 * * * /usr/bin/echo "Hello"' | crontab -u david -
```

The five scheduling fields mean minute `30`, hour `5`, every day of the month, every month, and every day of the week. Cron sends standard output to the user's local mail unless output is redirected.

## QUESTION NO: 7 - (SIMULATION)

**Install the Redhat Linux RHEL 4 through NFS. Where your Server is server1.example.com having IP 172.24.254.254 and shared /var/ftp/pub. The size of the partitions are listed below:**

**/?1048**

**/home?1028**

**/boot? 512**

**/var? 1028**

**/usr?2048**

**Swap->1.5 of RAM Size**

**/archive?configure the RAID Level 0 of remaining all free space.**

**After completing the installation through NFS solve the following questions. There are two networks 172.24.0.0/16 and 172.25.0.0/16. As well as there are two domains example.com on 172.24.0.0/16 network and my133t.org on 172.25.0.0/16 network. Your system is based on example.com domain.**

**ANSWER: See the explanation for the answer**

**Explanation:**

**Install using the NFS installation source:**

1. Boot the RHEL 4 installation media and at the boot prompt enter `linux askmethod`.
2. Configure the installer network interface with an unused address on the `172.24.0.0/16` network (netmask `255.255.0.0`). The installation server is `172.24.254.254`.
3. Select **NFS image** as the installation method.
4. Enter the following NFS source values:

```
NFS server: 172.24.254.254
Red Hat directory: /var/ftp/pub
```

5. Set the installed system hostname to a name in the `example.com` domain, for example `client.example.com`. Do not use the `my133t.org` domain for this system.

**Use custom disk partitioning** and create the required filesystems (sizes are in MiB/MB as displayed by Anaconda):

| Mount point | Size | Filesystem/notes |
|---|---|---|
| `/boot` | 512 | Standard non-RAID partition; ext3. |
| `/` | 1048 | ext3. |
| `/home` | 1028 | ext3. |
| `/var` | 1028 | ext3. |
| `/usr` | 2048 | ext3. |
| swap | 1.5 × installed RAM | Create a swap partition sized at one and one-half times the physical memory reported by the installer. |
| `/archive` | All remaining available disk space | Create as an **RAID 0** device and format it (normally ext3). |

**RAID 0 requirement:** create at least two RAID-member partitions from all space remaining after the fixed partitions and swap. Mark each as *software RAID*, then create an **RAID device** with RAID level **0 (stripe)**, using those members. Format that RAID device and assign its mount point as `/archive`. RAID 0 provides capacity/performance only; it has no redundancy.

After confirming the partition layout, allow Anaconda to format the filesystems and complete the package installation. The key NFS values are `172.24.254.254:/var/ftp/pub`, and the system belongs to `example.com` on the `172.24.0.0/16` network.

## QUESTION NO: 8 - (SIMULATION)

**You are the administrator of example.com domain. Configure to deny local login to all normal users on your domain server. As well as allow to root login only on First Terminal.**

## ANSWER: See the explanation for the answer

**Explanation:**

Enable PAM access control for the local console `login` service and then allow only root on `tty1`.

```
# Add this to /etc/pam.d/login (in the auth section)
auth       required     pam_access.so
```

Then configure `/etc/security/access.conf` with these rules, in this exact order:

```
+:root:tty1
-:ALL:LOCAL
```

The first rule permits `root` only when logging in on the first virtual terminal, `tty1`.
The second rule denies every account (including root on any other terminal) for local logins.
Network-based logins are not affected by the `LOCAL` rule.

The rule order is essential: PAM access rules are evaluated from top to bottom, and the first matching rule is used.

## QUESTION NO: 9 - (SIMULATION)

SIMULATION

Make Secondary belongs the both users on sysadmin group.

## ANSWER: See the explanation for the answer

**Explanation:**

Add both users to `sysadmin` as a supplementary (secondary) group without removing any existing supplementary-group memberships:

```
usermod -aG sysadmin john
usermod -aG sysadmin jane
```

The `-G` option sets supplementary groups, and `-a` appends the specified group. Omitting `-a` could replace the users' current supplementary groups.

Verify the result with:

```
id john
id jane
# or
getent group sysadmin
```

The users may need to log out and log back in before their new group membership is active in an existing login session.

## QUESTION NO: 10 - (SIMULATION)

**Make Successfully Resolve to server1.example.com where DNS Server is 192.168.0.254.**

## ANSWER: See the explanation for the answer

**Explanation:**

Configure the system DNS resolver to use `192.168.0.254` as its name server.

```
# printf 'nameserver 192.168.0.254\n' > /etc/resolv.conf
# host server1.example.com
```

The lookup should return an address for `server1.example.com`.

If NetworkManager manages the connection, make the setting persistent through its active connection profile instead of relying solely on a manually edited `/etc/resolv.conf`:

```
# nmcli connection show --active
# nmcli connection modify <connection-name> ipv4.dns 192.168.0.254 ipv4.ignore-auto-dns yes
# nmcli connection up <connection-name>
# host server1.example.com
```

## QUESTION NO: 11 - (SIMULATION)

**There is a NFS server 192.168.0.254 and all required packages are dumped in /var/ftp/pub of that server and the /var/ftp/pub directory is shared. Install the Redhat Enterprise Linux 5 by creating following partitions:**

**/1000**

**/boot200**

**/home1000**

**/var1000**

**/usr4000**

**swap2X256 (RAM SIZE)**

## ANSWER: See the explanation for the answer

**Explanation:**

Boot the RHEL 5 installation media and start an NFS-based installation. At the installer boot prompt, use `linux askmethod` if the installer does not automatically ask for an installation source.

Configure the network interface with an address that can reach `192.168.0.254`.
Select **NFS** as the installation method.
Enter NFS server: `192.168.0.254`
Enter exported directory: `/var/ftp/pub`

In Disk Druid/manual partitioning, create the following filesystems (sizes are in MB). Use `ext3`, the normal default filesystem for RHEL 5, for every non-swap partition:

| Mount point | Size | Filesystem type |
|---|---|---|
| `/boot` | 200 MB | ext3 |
| `/` | 1000 MB | ext3 |
| `/home` | 1000 MB | ext3 |
| `/var` | 1000 MB | ext3 |
| `/usr` | 4000 MB | ext3 |
| swap | 512 MB | swap |

The system has 256 MB RAM, so swap must be twice RAM: `2 × 256 MB = 512 MB`. Make `/boot` a standard disk partition (preferably primary); the other filesystems may be primary or logical partitions as appropriate for the disk layout. Continue the installation using the packages available from the NFS export.

## QUESTION NO: 12 - (SIMULATION)

**Create the directory /archive and group owner should be the sysuser group.**

## ANSWER: See the explanation for the answer

**Explanation:**

Create the directory and set its group ownership to `sysuser`:

```
mkdir -p /archive
chgrp sysuser /archive
```

Alternatively, in one command:

```
mkdir -p /archive && chown :sysuser /archive
```

Verify the group owner with:

```
ls -ld /archive
```

The group column should show `sysuser`.

## QUESTION NO: 13 - (SIMULATION)

**You have a ftp server having IP address 192.168.0.254. Using iptables, allow the ftp connection only from the internal network where internal network is 192.168.0.0/24.**

## ANSWER: See the explanation for the answer

**Explanation:**

Add an INPUT rule that permits FTP control connections (TCP/21) only when the source is the internal subnet, followed by a rule that denies TCP/21 from every other source.

```
iptables -I INPUT 1 -p tcp -s 192.168.0.0/24 -d 192.168.0.254 --dport 21 -j ACCEPT
iptables -I INPUT 2 -p tcp -d 192.168.0.254 --dport 21 -j DROP
```

The rule order is important: the subnet-specific `ACCEPT` must be evaluated before the general `DROP`.

If stateful FTP data connections are required, ensure established/related traffic is allowed (normally this rule already exists):

```
iptables -I INPUT 1 -m conntrack --ctstate ESTABLISHED,RELATED -j ACCEPT
```

Save the firewall rules using the persistence method appropriate for the installed Red Hat release, for example:

```
service iptables save
```

## QUESTION NO: 14 - (SIMULATION)

SIMULATION

There are mixed lots of System running on Linux and Windows OS. Some users are working on Windows Operating System. You want to make available /data directory to samba users only from 192.168.0.0/24 network. Configure the samba server.

## ANSWER: See the explanation for the answer

**Explanation:**

Configure an authenticated Samba share for `/data` and restrict Samba client connections to `192.168.0.0/24`. Do not enable guest access.

```
# Install the service if necessary
dnf install -y samba samba-client

# Create a group for users permitted to use this share
# (add the required existing Linux users to this group)
groupadd -f smbusers
usermod -aG smbusers USERNAME
smbpasswd -a USERNAME

# Prepare the shared directory
mkdir -p /data
chgrp smbusers /data
chmod 2770 /data

# Label the directory for Samba when SELinux is enforcing
semanage fcontext -a -t samba_share_t '/data(/.*)?'
restorecon -Rv /data
```

Add the following to `/etc/samba/smb.conf` (or set the equivalent options in the existing sections):

```
[global]
        security = user
        hosts allow = 192.168.0.0/24
        hosts deny = 0.0.0.0/0

[data]
        path = /data
        browseable = yes
        read only = no
        writable = yes
        guest ok = no
        valid users = @smbusers
```

`hosts allow` restricts access based on the Samba client IP address. `guest ok = no` requires Samba authentication, and `valid users = @smbusers` limits the share to members of the Linux group `smbusers`. Create each Samba login with `smbpasswd -a USERNAME`.

```
# Validate and start persistently
testparm
systemctl enable --now smb nmb

# Permit SMB through the firewall, if firewalld is used
firewall-cmd --permanent --add-service=samba
firewall-cmd --reload
```

On older RHEL systems, use `service smb restart` and `chkconfig smb on` instead of the `systemctl` commands.

## QUESTION NO: 15 - (SIMULATION)

SIMULATION

Create the directory /data and group owner should be the sysadmin group.

## ANSWER: See the explanation for the answer

**Explanation:**

Create the directory and assign its group ownership to `sysadmin`:

```
mkdir /data
```

```
chgrp sysadmin /data
```

Alternatively, perform both actions with:

```
mkdir /data
chown :sysadmin /data
```

Verify the group owner:

```
ls -ld /data
```

The listing should show `sysadmin` in the group column, for example:

```
drwxr-xr-x. 2 root sysadmin ... /data
```

## QUESTION NO: 16 - (SIMULATION)

**Configure to allow the pop3 and imap connection from your domain example.com**

**and my133t.org domain.**

## ANSWER: See the explanation for the answer

**Explanation:**

Configure Dovecot to permit POP3/IMAP connections only from clients in the `example.com` and `my133t.org` domains. Use Dovecot's `login_access_sockets` setting with DNS lookups for domain matching.

```
cat > /etc/dovecot/conf.d/99-login-access.conf <<'EOF'
login_access_sockets = tcpwrap
login_access_sockets = tcp:example.com tcp:my133t.org
EOF
```

More directly, on systems/configurations where the setting is placed in `/etc/dovecot/dovecot.conf`, add:

```
login_access_sockets = tcp:example.com tcp:my133t.org
```

Then validate and restart Dovecot:

```
dovecot -n
systemctl restart dovecot
systemctl enable dovecot
```

If the exam environment instead expects TCP Wrappers access control, add the permitted domains to `/etc/hosts.allow` and deny all other Dovecot clients in `/etc/hosts.deny`:

```
echo 'dovecot: .example.com .my133t.org' >> /etc/hosts.allow
echo 'dovecot: ALL' >> /etc/hosts.deny
systemctl restart dovecot
```

The intended result is that POP3/IMAP service access is accepted for hosts whose reverse DNS names belong to `example.com` or `my133t.org`.

## QUESTION NO: 17 - (SIMULATION)

SIMULATION

You have a domain in your LAN named example.com. Allow the FTP connection only from local domain.

## ANSWER: See the explanation for the answer

**Explanation:**

Configure TCP Wrappers so that `vsftpd` is denied to every client except hosts in `example.com`.

```
# vi /etc/hosts.deny
vsftpd: ALL EXCEPT .example.com
```

The leading dot means all hosts in that domain, such as `client1.example.com`. Clients outside the domain match `ALL` and are denied, while clients in `.example.com` do not match the deny rule and are permitted.

Ensure there is no conflicting `vsftpd` denial in `/etc/hosts.allow` or other access-control configuration.

## QUESTION NO: 18 - (SIMULATION)

SIMULATION

Configure the DHCP server by matching the following conditions:

1. Subnet and netmask should be 192.168.0.0 255.255.255.0

2. Gateway Should be 192.168.0.254

3. DNS Sever Should be 192.168.0.254

4. Domain Name should be example.com

5. Range from 192.168.0.10-50

## ANSWER: See the explanation for the answer

**Explanation:**

Edit the DHCP server configuration file, normally `/etc/dhcp/dhcpd.conf`, and configure the following scope:

```
ddns-update-style none;

default-lease-time 21600;
max-lease-time 43200;

option domain-name "example.com";
option domain-name-servers 192.168.0.254;
option routers 192.168.0.254;

subnet 192.168.0.0 netmask 255.255.255.0 {
    range 192.168.0.10 192.168.0.50;
}
```

The important correction is that the gateway/router option must be `192.168.0.254`, not `192.168.0.1`.

Validate the configuration and start/enable the service:

```
dhcpd -t -cf /etc/dhcp/dhcpd.conf
systemctl enable --now dhcpd
```

If the service was already running, apply the changed configuration with:

```
systemctl restart dhcpd
```

## QUESTION NO: 19 - (SIMULATION)

**One New Kernel is released named kernel-.2.6.9-11. Kernel is available on** ftp://server1.example.com**/pub/updates directory for anonymous. Install the Kernel and make the kernel-2.6.9-5 default.**

**ANSWER: See the explanation for the answer**

**Explanation:**

Download and install the new kernel *without upgrading/removing* the existing kernel. Kernel packages must be installed with `-i`, not `-U`, because kernel-2.6.9-5 must remain available.

```
# cd /tmp
# wget ftp://server1.example.com/pub/updates/kernel-2.6.9-11*.rpm
# rpm -ivh kernel-2.6.9-11*.rpm
```

(Use the exact RPM filename shown in the FTP directory, including its architecture, for example `kernel-2.6.9-11.i686.rpm`.)

Then edit the legacy GRUB configuration:

```
# vi /boot/grub/grub.conf
```

Locate the `title` entry that boots `kernel-2.6.9-5`. Set the global `default=` value to that entry's zero-based position. For the usual result after installing 2.6.9-11, the new kernel is entry 0 and 2.6.9-5 is entry 1, so set:

```
default=1
```

For example:

```
default=1
timeout=5
...
title Red Hat Enterprise Linux (...2.6.9-11)
        ...
title Red Hat Enterprise Linux (...2.6.9-5)
        ...
```

Confirm the installed kernels and GRUB selection:

```
# rpm -q kernel
# grep -E '^(default|title)' /boot/grub/grub.conf
```

Reboot when required; GRUB will boot kernel-2.6.9-5 by default.

**QUESTION NO: 20 - (SIMULATION)**

**All mails to my133t.org should get by marion user.**

**ANSWER: See the explanation for the answer**

**Explanation:**

Configure a Postfix virtual alias catch-all for the domain so that every recipient at `my133t.org` is delivered to local user `marion`.

```
echo '@my133t.org marion' >> /etc/postfix/virtual
postmap /etc/postfix/virtual
postconf -e 'virtual_alias_maps = hash:/etc/postfix/virtual'
systemctl reload postfix
```

The essential mapping is:

```
@my133t.org    marion
```

This catches addresses such as `anything@my133t.org` and rewrites their delivery recipient to `marion`. Ensure the local account exists if needed:

```
id marion || useradd marion
```

## QUESTION NO: 21 - (SIMULATION)

**Make on /data that only the user owner and group owner member can fully access.**

## ANSWER: See the explanation for the answer

**Explanation:**

Set restrictive permissions on `/data` so that the owning user and owning group have full access, while all other users have no access:

```
chmod 770 /data
```

This results in permissions `rwxrwx---`:

Owner: read, write, and enter/access the directory.
Group: read, write, and enter/access the directory.
Others: no permissions.

Verify with:

```
ls -ld /data
```

The mode should display as `drwxrwx---` (assuming `/data` is a directory).

## QUESTION NO: 22 - (SIMULATION)

**There are mixed lots of System running on Linux and Windows OS. Some users are working on Windows Operating System. There is a /data directory on linux server should make available on windows to only user1 and user2 users with full access. Configure to make available.**

## ANSWER: See the explanation for the answer

**Explanation:**

Configure a Samba share for `/data`, restrict it to `user1` and `user2`, and allow both users read/write access.

```
# Install Samba packages
sudo dnf install -y samba samba-client

# Ensure the directory exists and is writable by the permitted users
sudo mkdir -p /data
sudo chgrp user1 /data
sudo chmod 0770 /data
# If user2 is not already in user1's group, create/use a common group instead:
# sudo groupadd datausers
# sudo usermod -aG datausers user1
# sudo usermod -aG datausers user2
# sudo chgrp datausers /data
# sudo chmod 2770 /data

# Create Samba passwords (the Linux accounts must already exist)
sudo smbpasswd -a user1
sudo smbpasswd -a user2
sudo smbpasswd -e user1
sudo smbpasswd -e user2
```

Add the following share definition to `/etc/samba/smb.conf`:

```
[data]
    comment = Shared data for user1 and user2
    path = /data
    browseable = yes
```

```
    read only = no
    writable = yes
    valid users = user1 user2
    force group = datausers
    create mask = 0660
    directory mask = 0770
```

If a common group was not used, omit `force group = datausers`. The recommended configuration is to use the shared `datausers` group so that files created by either user remain editable by the other user.

```
# Validate Samba configuration
sudo testparm

# Apply SELinux labels and permit Samba read/write sharing
sudo semanage fcontext -a -t samba_share_t '/data(/.*)?'
sudo restorecon -Rv /data
sudo setsebool -P samba_export_all_rw on

# Start Samba and enable it at boot
sudo systemctl enable --now smb nmb

# Permit Windows SMB access through the firewall
sudo firewall-cmd --permanent --add-service=samba
sudo firewall-cmd --reload
```

From Windows, open `\\<linux-server-name-or-IP>\data` and authenticate as either `user1` or `user2` using the Samba password configured with `smbpasswd`. The `valid users` directive prevents all other Samba users from accessing this share.

## QUESTION NO: 23 - (SIMULATION)

**Configure the caching only-name server for example.com where DNS server is 192.100.0.254.**

## ANSWER: See the explanation for the answer

**Explanation:**

Install and configure `bind` as a caching-only resolver. Do not create a master zone or zone files for `example.com`; cached answers must be obtained from the upstream DNS server `192.100.0.254`.

```
yum -y install bind
```

Edit `/etc/named.conf` and ensure its `options` block contains the following (retain the standard `directory` setting if it already exists):

```
options {
        directory "/var/named";
        listen-on port 53 { any; };
        listen-on-v6 { none; };
        allow-query { any; };
        recursion yes;

        forward only;
        forwarders { 192.100.0.254; };
};
```

If the requirement is to forward *only* the `example.com` namespace, rather than all DNS queries, use this forward zone instead of global `forward only`/`forwarders` directives:

```
zone "example.com" IN {
        type forward;
        forward only;
        forwarders { 192.100.0.254; };
```

```
};
```

Validate and start the service:

```
named-checkconf
systemctl enable --now named
firewall-cmd --permanent --add-service=dns
firewall-cmd --reload
```

Test that the server can cache and resolve names from the domain:

```
dig @127.0.0.1 example.com
```

The first lookup is forwarded to `192.100.0.254`; subsequent lookups are answered from the local BIND cache until the DNS TTL expires.

## QUESTION NO: 24 - (SIMULATION)

**/storage directory on linux server should make available on windows system that jeff user should able to access on read only mode within example.com domain.**

## ANSWER: See the explanation for the answer

**Explanation:**

Configure Samba to export `/storage` as a share that only `jeff` can connect to and which is read-only. The Windows client can then use `\\<server>.example.com\storage`.

```
dnf install -y samba samba-client
# Ensure jeff is a local account (if it does not already exist)
id jeff || useradd jeff
# Set the Samba password for the Windows/Samba login
smbpasswd -a jeff

# Allow jeff to traverse/read the exported content at the filesystem level
setfacl -R -m u:jeff:rX /storage

# Permit Samba to export this directory under SELinux
semanage fcontext -a -t samba_share_t '/storage(/.*)?'
restorecon -Rv /storage
```

Add the following to `/etc/samba/smb.conf` (or merge it with the existing configuration):

```
[global]
        workgroup = EXAMPLE
        security = user

[storage]
        path = /storage
        comment = Read-only storage share
        browseable = yes
        read only = yes
        valid users = jeff
```

`EXAMPLE` is the NetBIOS workgroup name corresponding to the `example.com` environment. The important access controls are `valid users = jeff` and `read only = yes`; the latter prevents SMB writes even if filesystem permissions would otherwise permit them.

```
testparm
systemctl enable --now smb nmb
firewall-cmd --permanent --add-service=samba
firewall-cmd --reload
```

From Windows, connect as `EXAMPLE\jeff` (or the appropriate `jeff` Samba credential) to `\\server.example.com\storage`. Jeff will be able to read/list the share but cannot create, modify, or delete files through Samba.

## QUESTION NO: 25 - (SIMULATION)

SIMULATION

Raw printer named printerx where x is your station number is installed and shared on server1.example.com. Install the shared printer on your PC to connect shared printer using IPP Protocols. Your server is 192.168.0.254.

## ANSWER: See the explanation for the answer

**Explanation:**

Add a local CUPS queue that uses the remote shared raw queue through IPP.

1. If `server1.example.com` does not already resolve, add the server address to name resolution:

   ```
   echo '192.168.0.254 server1.example.com server1' >> /etc/hosts
   ```

2. Open the local CUPS administration page at `http://localhost:631`, choose **Administration** / **Add Printer**, and authenticate as root if requested.
3. Create a local queue (normally use `printerx`, where `x` is the station number; a different local queue name is also acceptable if required by the lab).
4. Select **Internet Printing Protocol (IPP)** as the connection type.
5. Use this device URI, substituting the station number:

   ```
   ipp://server1.example.com/printers/printerx
   ```

   If DNS is unavailable and the server IP must be used directly, use:

   ```
   ipp://192.168.0.254/printers/printerx
   ```

6. Select the **Raw** / **Raw Queue** printer model and finish the configuration. Enable the queue and print a test page if requested.

Equivalent command-line configuration:

```
echo '192.168.0.254 server1.example.com server1' >> /etc/hosts
lpadmin -p printerx -E \
  -v ipp://server1.example.com/printers/printerx \
  -m raw
lpoptions -d printerx
lpstat -v printerx
```

The important settings are the IPP URI `ipp://server1.example.com/printers/printerx` and the **Raw** model, because the shared printer on `server1` is a raw printer.

## QUESTION NO: 26 - (SIMULATION)

SIMULATION

There are mixed lots of System running on Linux and Windows OS. Some users are working on Windows Operating System. Your printer is connected on linux server. You want to share the printer-using samba so that users working on windows also can print. Configure the samba server to share printer.

## ANSWER: See the explanation for the answer

**Explanation:**

Configure CUPS first and ensure that the local printer can print successfully. Then export the installed CUPS printers through Samba.

```
dnf install -y samba samba-client cups
systemctl enable --now cups smb nmb
mkdir -p /var/spool/samba
chmod 1777 /var/spool/samba
restorecon -RFv /var/spool/samba
```

Add (or ensure) the following configuration in `/etc/samba/smb.conf`:

```
[global]
        workgroup = WORKGROUP
        security = user
        printing = cups
        printcap name = cups
        load printers = yes

[printers]
        comment = All Printers
        path = /var/spool/samba
        browsable = yes
        printable = yes
        writable = no
        guest ok = no
```

The `[printers]` section is a special Samba share: every printer known to CUPS is made available to Windows clients. `printable = yes` makes this a printer share, and `/var/spool/samba` is used for temporary print-job spooling.

Validate and activate the configuration:

```
testparm
systemctl restart smb nmb
firewall-cmd --permanent --add-service=samba
firewall-cmd --reload
```

If SELinux has a non-default label on the spool directory, restore its expected label before restarting Samba:

```
restorecon -RFv /var/spool/samba
```

For authenticated access, create Samba credentials for each allowed Linux user:

```
smbpasswd -a username
```

Windows users can then browse to `\\<samba-server>\` and connect to the desired printer, or add it as `\\<samba-server>\<CUPS-printer-name>`.

## QUESTION NO: 27 - (SIMULATION)

**One NIS Domain named rhce.com is configured in your lab, server is**

**192.168.0.254. rhce100, rhce200,rhce300 user are created on domain server.**

**Make your system as a member of rhce.com domain. Make sure that when nis user login in your system home directory should get by them. Home directory is separately shared on server eg /home/stationx/ where x is you station number.**

## ANSWER: See the explanation for the answer

**Explanation:**

Configure the machine as an NIS client for `rhce.com`, then mount this workstation's exported home-directory tree from the NIS/NFS server. Replace `X` below with the assigned station number (for example, use `/home/station1` for station 1).

```
# Configure NIS client authentication
cat > /etc/yp.conf <<'EOF'
domain rhce.com server 192.168.0.254
EOF

authconfig --enablenis --nisdomain=rhce.com --nisserver=192.168.0.254 --enablemkhomedir --
update

# Start NIS binding services and make them persistent
systemctl enable --now rpcbind ypbind
```

On systems using the older SysV service commands, the equivalent is:

```
service rpcbind start
service ypbind start
chkconfig rpcbind on
chkconfig ypbind on
```

Ensure that NIS is used for account lookups. Normally `authconfig` has already made these changes; the relevant entries in `/etc/nsswitch.conf` must include NIS, for example:

```
passwd:     files nis
shadow:     files nis
group:      files nis
netgroup:   nis
```

Mount the station-specific NFS export at `/home` so that an NIS passwd entry such as `/home/rhce100` resolves to the correct shared home directory:

```
mkdir -p /home
printf '%s\n' '192.168.0.254:/home/stationX /home nfs defaults,_netdev 0 0' >> /etc/fstab
mount -a
```

For example, on station 2 the fstab line is:

```
192.168.0.254:/home/station2 /home nfs defaults,_netdev 0 0
```

`--enablemkhomedir` enables PAM home-directory creation at first login. If the NFS server already contains the users' directories, it simply allows those existing shared homes to be used after the NFS mount.

Validate the configuration:

```
ypwhich
getent passwd rhce100
mount | grep ' on /home '
su - rhce100
pwd
```

`ypwhich` should report `192.168.0.254`, `getent` should return the NIS user, and after `su - rhce100` the working directory should be that user's shared home under `/home`.

## QUESTION NO: 28 - (SIMULATION)

SIMULATION

Port 8080

Configure the squid server to allow the Local Domain and deny to cracker.org domain.

## ANSWER: See the explanation for the answer

**Explanation:**

Edit `/etc/squid/squid.conf` and configure Squid to listen on port `8080`. Use destination-domain ACLs; do not use source-network ACLs for this requirement.

```
http_port 8080

# Replace localdomain with the DNS domain configured for this host/network.
# For example, if /etc/resolv.conf contains "search example.com":
acl local_domain dstdomain .example.com
acl cracker_domain dstdomain .cracker.org

# This denial must be before any general allow rule such as "allow localnet".
http_access deny cracker_domain
http_access allow local_domain
http_access deny all
```

The leading dot means the ACL also matches subdomains. Determine the actual local DNS suffix from `/etc/resolv.conf` (the `search` or `domain` line) or with `hostname -d`, and use it in place of `.example.com`. For a local domain named `localdomain`, use:

```
acl local_domain dstdomain .localdomain
```

If the default configuration contains an earlier broad rule such as `http_access allow localnet`, place `http_access deny cracker_domain` before it; Squid evaluates `http_access` rules in order.

Validate, start, and enable the service:

```
squid -k parse
systemctl enable --now squid
```

Squid will now accept proxy connections on TCP port `8080`, permit requests to the local DNS domain, and explicitly deny requests to `cracker.org` and its subdomains.

## QUESTION NO: 29 - (SIMULATION)

SIMULATION

Create the partition having 100MB size and mount it on /mnt/neo

## ANSWER: See the explanation for the answer

**Explanation:**

Create a new 100 MiB partition on the available unused disk, format it, mount it at `/mnt/neo`, and make the mount persistent. For example, if the unused disk is `/dev/sdb` and the new partition becomes `/dev/sdb1`:

```
# fdisk /dev/sdb
# In fdisk:
n                 # new partition
p                 # primary partition (or accept the appropriate default)
1                 # partition number
<Enter>           # default first sector
+100M             # size
w                 # write changes

# partprobe /dev/sdb
# mkfs.xfs -f /dev/sdb1
# mkdir -p /mnt/neo
# mount /dev/sdb1 /mnt/neo
# blkid /dev/sdb1
```

Add the UUID returned by `blkid` to `/etc/fstab`:

```
UUID=<uuid-of-/dev/sdb1>  /mnt/neo  xfs  defaults  0  0
```

Validate the persistent configuration:

```
# umount /mnt/neo
# mount -a
# findmnt /mnt/neo
```

If the system or exam requires an ext filesystem instead, use `mkfs.ext4 -F /dev/sdb1` (and specify `ext4` in `/etc/fstab`) rather than XFS. Replace `/dev/sdb` and `/dev/sdb1` with the actual unused disk and resulting partition shown by `lsblk`.

## QUESTION NO: 30 - (SIMULATION)

**If any mail coming from outside of the local LAN block all mails.**

## ANSWER: See the explanation for the answer

**Explanation:**

Configure Postfix to permit SMTP connections only from the loopback address and the local LAN, and reject every other SMTP client. Edit `/etc/postfix/main.cf` (replace `192.168.1.0/24` with the actual local-LAN subnet):

```
mynetworks = 127.0.0.0/8, [::1]/128, 192.168.1.0/24
smtpd_client_restrictions = permit_mynetworks, reject
```

Apply and verify the configuration:

```
postfix check
systemctl restart postfix
systemctl enable postfix
```

`permit_mynetworks` accepts SMTP clients from the configured LAN, while the final `reject` blocks mail submission attempts from any address outside that LAN. Do not use a broad network such as `0.0.0.0/0` in `mynetworks`.

## QUESTION NO: 31 - (SIMULATION)

**There are mixed lots of System running on Linux and Windows OS. Some users are working on Windows Operating System. There is a /data directory on linux server should make available on windows to user1 and user2 users on read and write mode and read only to other samba users.**

## ANSWER: See the explanation for the answer

**Explanation:**

Configure a Samba share for `/data` as read-only by default, and grant write access only to `user1` and `user2`.

```
# Install and start Samba
sudo dnf install -y samba samba-client policycoreutils-python-utils
sudo systemctl enable --now smb

# Create Samba access groups
sudo groupadd samba_rw
sudo groupadd samba_ro
sudo usermod -aG samba_rw user1
sudo usermod -aG samba_rw user2

# Add every other allowed Samba user to samba_ro
# Example:
# sudo usermod -aG samba_ro user3

# Set Samba passwords (repeat for each Samba user)
```

```
sudo smbpasswd -a user1
sudo smbpasswd -a user2

# Allow the two groups to access the directory at the Linux filesystem level
sudo chgrp samba_rw /data
sudo chmod 2770 /data
sudo setfacl -m g:samba_rw:rwx,g:samba_ro:rx /data
sudo setfacl -d -m g:samba_rw:rwx,g:samba_ro:rx /data

# Label the directory for Samba when SELinux is enforcing
sudo semanage fcontext -a -t samba_share_t '/data(/.*)?'
sudo restorecon -Rv /data
sudo setsebool -P samba_export_all_rw on

# Permit Samba through the firewall
sudo firewall-cmd --permanent --add-service=samba
sudo firewall-cmd --reload
```

Add the following share definition to `/etc/samba/smb.conf`:

```
[data]
        comment = Shared Data
        path = /data
        browseable = yes
        valid users = @samba_rw @samba_ro
        read only = yes
        write list = @samba_rw
        create mask = 0660
        directory mask = 2770
        force group = samba_rw
```

Validate and reload Samba:

```
sudo testparm
sudo systemctl restart smb
```

**Result:** all users placed in `samba_ro` can browse and read `/data`, while `user1` and `user2`, through membership in `samba_rw`, can create, modify, and delete files. Windows clients can connect using `\\<samba-server>\data`.

## QUESTION NO: 32 - (SIMULATION)

**Your Local Domain is example.com. Configure the send mail server for you local LAN by following these conditions.**

**i. Any mail going from Local LAN should be masquerade to example.com**

**ii. Any incoming mail for** info@example.com **virtual address should be mapped to admin@example.com**

**iii. All outgoing mail should be send via smtp.abc.com mail server.**

## ANSWER: See the explanation for the answer

**Explanation:**

Configure Postfix so that `example.com` is the local/masquerade domain, use a virtual alias map for the `info` address, and relay all non-local mail through `smtp.abc.com`.

Edit `/etc/postfix/main.cf` and ensure it contains the following settings (preserve any existing required settings):

```
mydomain = example.com
myorigin = $mydomain
mydestination = $myhostname, localhost.$mydomain, localhost, $mydomain
```

```
# Masquerade mail from hosts in the local example.com domain
masquerade_domains = $mydomain

# Virtual address translation
virtual_alias_maps = hash:/etc/postfix/virtual

# Send all outgoing/non-local mail through this SMTP relay
relayhost = [smtp.abc.com]
```

Create `/etc/postfix/virtual` with this virtual alias:

```
info@example.com    admin@example.com
```

Build the lookup database and reload Postfix:

```
postmap /etc/postfix/virtual
systemctl reload postfix
```

The brackets in `relayhost = [smtp.abc.com]` make Postfix connect directly to that SMTP host instead of performing MX-record lookup. Mail originating as addresses under the local `example.com` domain is masqueraded as `@example.com`, incoming mail for `info@example.com` is delivered to `admin@example.com`, and external delivery is relayed through `smtp.abc.com`.

## QUESTION NO: 33 - (SIMULATION)

**You Completely Install the Redhat Enterprise Linux 5 on your System. While start the system, it's giving error to load X window System. How will you fix that problem and make boot successfully run X Window System.**

## ANSWER: See the explanation for the answer

**Explanation:**

Boot to a text console, correct the X configuration, test it, and then return the system to graphical runlevel.

1. At the GRUB boot menu, edit the kernel line and append `3` (or use rescue/single-user mode if necessary). This boots into runlevel 3 without trying to start the display manager.
2. Log in as `root`.
3. Re-create/configure the X server configuration for the installed video adapter and monitor:

```
system-config-display --reconfig
```

Select the correct video driver, monitor, resolution, and color depth. On RHEL 5 this writes the X configuration, normally `/etc/X11/xorg.conf`.

If the configuration utility cannot be used graphically, a basic configuration can also be generated with:

```
mv /etc/X11/xorg.conf /etc/X11/xorg.conf.bad  # only if an invalid file exists
X -configure
cp /root/xorg.conf.new /etc/X11/xorg.conf
```

4. Test X before making it the default:

```
startx
```

If X fails, inspect the X error log and correct the indicated driver, monitor, or configuration issue:

```
grep -E '\(EE\)|\(WW\)' /var/log/Xorg.0.log
```

5. Ensure graphical mode is the default boot target. In RHEL 5, edit `/etc/inittab` and set:

```
id:5:initdefault:
```

Then reboot, or switch immediately to graphical runlevel:

```
init 5
```

The key fix is to boot in text mode first, generate a valid `/etc/X11/xorg.conf` using `system-config-display --reconfig`, verify with `startx`, and set the default runlevel to `5`.

## QUESTION NO: 34 - (SIMULATION)

**There is a NFS server 192.168.0.254 and all required packages are dumped in /var/ftp/pub of that server and the /var/ftp/pub directory is shared. Install the Redhat Enterprise Linux 5 by creating following partitions:**

**/ 1000**

**/boot 200**

**/home 1000**

**/var 1000**

**/usr 4000**

**swap 2X256 (RAM SIZE)**

## ANSWER: See the explanation for the answer

**Explanation:**

Boot the RHEL 5 installer and select a network installation. Because the server is specified as an NFS server and its shared installation tree is `/var/ftp/pub`, select **NFS directory** as the installation method and enter:

NFS server: `192.168.0.254`
RHEL installation directory: `/var/ftp/pub`

Configure network addressing as required to reach `192.168.0.254`. At the disk partitioning screen choose **Create custom layout** (manual partitioning). Create the following filesystems, using `ext3`, which is the normal RHEL 5 filesystem type:

| Mount point / type | Size |
|---|---|
| `/boot` | 200 MB |
| `/` | 1000 MB |
| `/home` | 1000 MB |
| `/var` | 1000 MB |
| `/usr` | 4000 MB |
| swap | 512 MB |

The system has 256 MB RAM, so the required swap size is `2 × 256 MB = 512 MB`. Ensure that the `/boot` filesystem is a normal filesystem partition (not swap), and set the listed mount points exactly before continuing with the installation.

## QUESTION NO: 35 - (SIMULATION)

SIMULATION

Eric user should able to write on Document root directory.

## ANSWER: See the explanation for the answer

**Explanation:**

Grant `eric` write access to the Apache DocumentRoot with a POSIX ACL. On the default Red Hat Apache configuration, the DocumentRoot is `/var/www/html`:

```
setfacl -m u:eric:rwx /var/www/html
getfacl /var/www/html
```

The resulting ACL should contain an entry similar to:

```
user:eric:rwx
```

If the site uses a nondefault DocumentRoot, first identify it and apply the ACL to that directory instead:

```
grep -R '^DocumentRoot' /etc/httpd/conf /etc/httpd/conf.d
setfacl -m u:eric:rwx /path/to/documentroot
```

If Eric must also automatically have access to files and directories created later below the DocumentRoot, add a default ACL:

```
setfacl -m d:u:eric:rwx /var/www/html
```

On current Red Hat filesystems, ACL support is normally enabled by default. If the DocumentRoot is on an older filesystem mounted without ACL support, enable the `acl` mount option in `/etc/fstab` for that filesystem and remount it before using `setfacl`.

## QUESTION NO: 36 - (SIMULATION)

**Create the user named jackie, curtin, david**

## ANSWER: See the explanation for the answer

**Explanation:**

Create the three local users with `useradd`:

```
useradd jackie
useradd curtin
useradd david
```

Verify that the accounts exist:

```
id jackie
id curtin
id david
```

No password, group membership, home-directory path, or shell is specified, so use the system defaults.

## QUESTION NO: 37 - (SIMULATION)

**Port 8080**

**Configure the squid server to allow the Local Domain and deny to my133t.org domain.**

## ANSWER: See the explanation for the answer

**Explanation:**

Edit `/etc/squid/squid.conf` and configure Squid to listen on TCP port `8080`. Add the access rules **before** the final `http_access deny all` rule:

```
http_port 8080

# Replace example.com with the DNS domain returned by: hostname -d
acl local_domain dstdomain .example.com
acl deny_my133t dstdomain .my133t.org
```

```
http_access deny deny_my133t
http_access allow local_domain
http_access deny all
```

The leading dot causes the ACL to match both the domain and its subdomains. The deny rule is placed first so that `my133t.org` is always rejected.

Validate and start the service:

```
squid -k parse
systemctl enable --now squid
```

If remote proxy clients must reach this server, permit the new listening port through the firewall:

```
firewall-cmd --permanent --add-port=8080/tcp
firewall-cmd --reload
```

## QUESTION NO: 38 - (SIMULATION)

**Some users home directory is shared from your system. Using showmount –e localhost command, the shared directory is not shown. Make access the shared users home directory.**

## ANSWER: See the explanation for the answer

**Explanation:**

Configure the users' home directory as an NFS export, reload the NFS export table, and ensure the NFS service is enabled. For example, export `/home` to the required client network (use the specified network instead of `*` if one is provided):

```
# vi /etc/exports
/home *(rw,sync)

# systemctl enable --now nfs-server
# exportfs -rav
```

If a firewall is in use, permit NFS access:

```
# firewall-cmd --permanent --add-service=nfs
# firewall-cmd --permanent --add-service=mountd
# firewall-cmd --permanent --add-service=rpc-bind
# firewall-cmd --reload
```

Verify that the home directory is now published:

```
# showmount -e localhost
Export list for localhost:
/home *
```

If an existing `/home` entry is already present in `/etc/exports`, do not duplicate it; correcting the entry and running `exportfs -rav` is sufficient.

## QUESTION NO: 39 - (SIMULATION)

SIMULATION

Install the Redhat Linux RHEL 4 through NFS. Where your Server is server1.example.com having IP 172.24.254.254 and shared /var/ftp/pub. The size of the partitions are listed below:

/ □ 1048

/home □ 1028

/boot □ 512

/var □ 1028

/usr □ 2048

Swap -> 1.5 of RAM Size

/archive □ configure the RAID Level 0 of remaining all free space.

After completing the installation through NFS solve the following questions. There are two networks 172.24.0.0/16 and 172.25.0.0/16. As well as there are two domains example.com on 172.24.0.0/16 network and my133t.org on 172.25.0.0/16 network. Your system is based on example.com domain.

**ANSWER: See the explanation for the answer**

**Explanation:**

Boot the RHEL 4 installer and start a network installation. At the boot prompt use:

```
linux askmethod
```

Select **NFS image** as the installation method. Use DHCP/dynamic network configuration if it is available in the lab. For the NFS installation source enter:

NFS server: `172.24.254.254` (or `server1.example.com` if name resolution is already working)
Red Hat Enterprise Linux directory: `/var/ftp/pub`

Continue through language, keyboard, and network screens, then choose custom/manual disk partitioning. Create these filesystems (the displayed units in this old installer are normally MB):

`/boot`: `512 MB`, ext3, preferably a standard primary partition (not RAID).
`/`: `1048 MB`, ext3.
`/home`: `1028 MB`, ext3.
`/var`: `1028 MB`, ext3.
`/usr`: `2048 MB`, ext3.
swap: create a swap partition sized to `1.5 × installed RAM`. For example, with 512 MB RAM create 768 MB swap; with 1 GB RAM create 1536 MB swap.

After allocating the required normal partitions and swap, use **all remaining disk space** to create two equal-size RAID member partitions. Mark both partitions as RAID members, then use the installer's **RAID** button to create an MD RAID device with:

RAID level: **RAID 0 (striped)**
Members: the two equal RAID member partitions
Mount point: `/archive`
Filesystem: ext3 (the normal RHEL 4 choice)

RAID 0 provides capacity and striping only; it has no fault tolerance. Ensure that the two RAID member partitions consume every remaining unallocated extent, so that the resulting RAID device mounted on `/archive` uses all free disk space.

Install the boot loader to the MBR, select the required timezone and root password, choose the required package groups for the lab, complete installation, and reboot. The installed system belongs to the `example.com` environment; the NFS source is reached from the `172.24.0.0/16` network via `172.24.254.254`.