

Network Address Translation (NAT)

Dynamic Host Configuration Protocol (DHCP)

OBJECTIVES

- How NAT (network address translation) works
- How DHCP (Dynamic Host Configuration Protocol) works
- How DHCP works together with NAT

CONTENTS

PRELAB 7		253
	Question Sheet for Prelab 7	255
LAB 7		257
PART 1	Network Address Translation (NAT)	258
PART 2	Dynamic Host Configuration Protocol (DHCP)	266
PART 3	Combining NAT and DHCP	272
	<i>Checklist Form for Lab 7</i>	275
	<i>Feedback Form for Lab 7</i>	277

Please note that you only need to do this part.

Pre-lab is optional

PRELAB 7

1. **Unix commands for NAT, DHCP:** Go to the online manual pages at

<http://www.tcpip-lab.net/links/manual.html>

and select the operating system version Red Hat Linux/i386 7.3. Read the manual pages of the following commands:

- a. `iptables`
- b. `dhcpd`
- c. `dhcpd.conf`
- d. `dhcp-options`
- e. `dhcpd.leases`

Read the manual page of `dhclient` at

<http://www.tcpip-lab.net/links/dhclient.html>

2. **Private IP addresses:** Read RFC 1918 on address allocation in private networks.

<http://www.tcpip-lab.net/links/rfc1918.html>

3. **Network address translation (NAT):** Read the tutorial on NAT at

<http://www.tcpip-lab.net/links/nat.html>

4. **Netfilter/iptables:** Read about *netfilter* and `iptables` at

<http://www.tcpip-lab.net/links/netfilter.html>

5. **Dynamic Host Configuration Protocol (DHCP):** Read RFC 2131 on DHCP at

<http://www.tcpip-lab.net/links/rfc2131.html>

QUESTION SHEET FOR PRELAB 7 **Optional**

Answer the questions in the space provided below each one. Use extra sheets of paper if needed and attach them to this document. Submit the answers to the prelab with your lab report.

Name (please print): _____

1. Explain why NAT is often mentioned as a solution to counteract the depletion of IP addresses on the global Internet. Which alternatives to NAT exist that address the scarcity of available IP addresses?

2. What does the following comment refer to: "NAT destroys the ability to do host-to-host communication over the Internet"?

3. Explain the following terms as used in the context of Network Address Translation:
 - a. Static NAT

 - b. Dynamic NAT

 - c. NAT with IP overload

 - d. Port address translations

 - e. IP masquerading

4. Refer to RFC 1918 and list the IP address blocks that are reserved for use in private networks. Why is there a need to specify IP addresses for private networks?

5. The utility *netfilter* and the command `iptables` provide support for NAT in Linux systems. Explain the relationship between the *netfilter* utility and the `iptables` command.

6. Describe the following terms that are used in the `iptables` command:
 - a. Chain

 - b. Postrouting

 - c. Prerouting

7. Consider a NAT device between a private and the public network. Suppose the private network uses addresses in the range 10.0.1.0–10.0.1.255, and suppose that the interface of the NAT device to the public network has IP address 128.143.136.80.
 - a. Write the `iptables` command so that the addresses in the private network are mapped to the public IP address 128.143.136.80.

 - b. Write an IOS command so that the addresses in the private network are mapped to the public IP address 128.143.136.80.

8. Explain the meaning of the *magic cookie* in the DHCP protocol.

9. If the command `dhcpcd` is issued (without arguments) on a Linux PC with multiple network interfaces, which network interfaces does the DHCP server listen on?

You are NOT required to finished Part 1 of lab 7.

LAB 7

Figure 7.1 shows two private networks that are connected to a public network. Each private network is connected to the public network by a NAT device, which is either a PC or a Cisco router. On each NAT device, IP forwarding must be enabled.

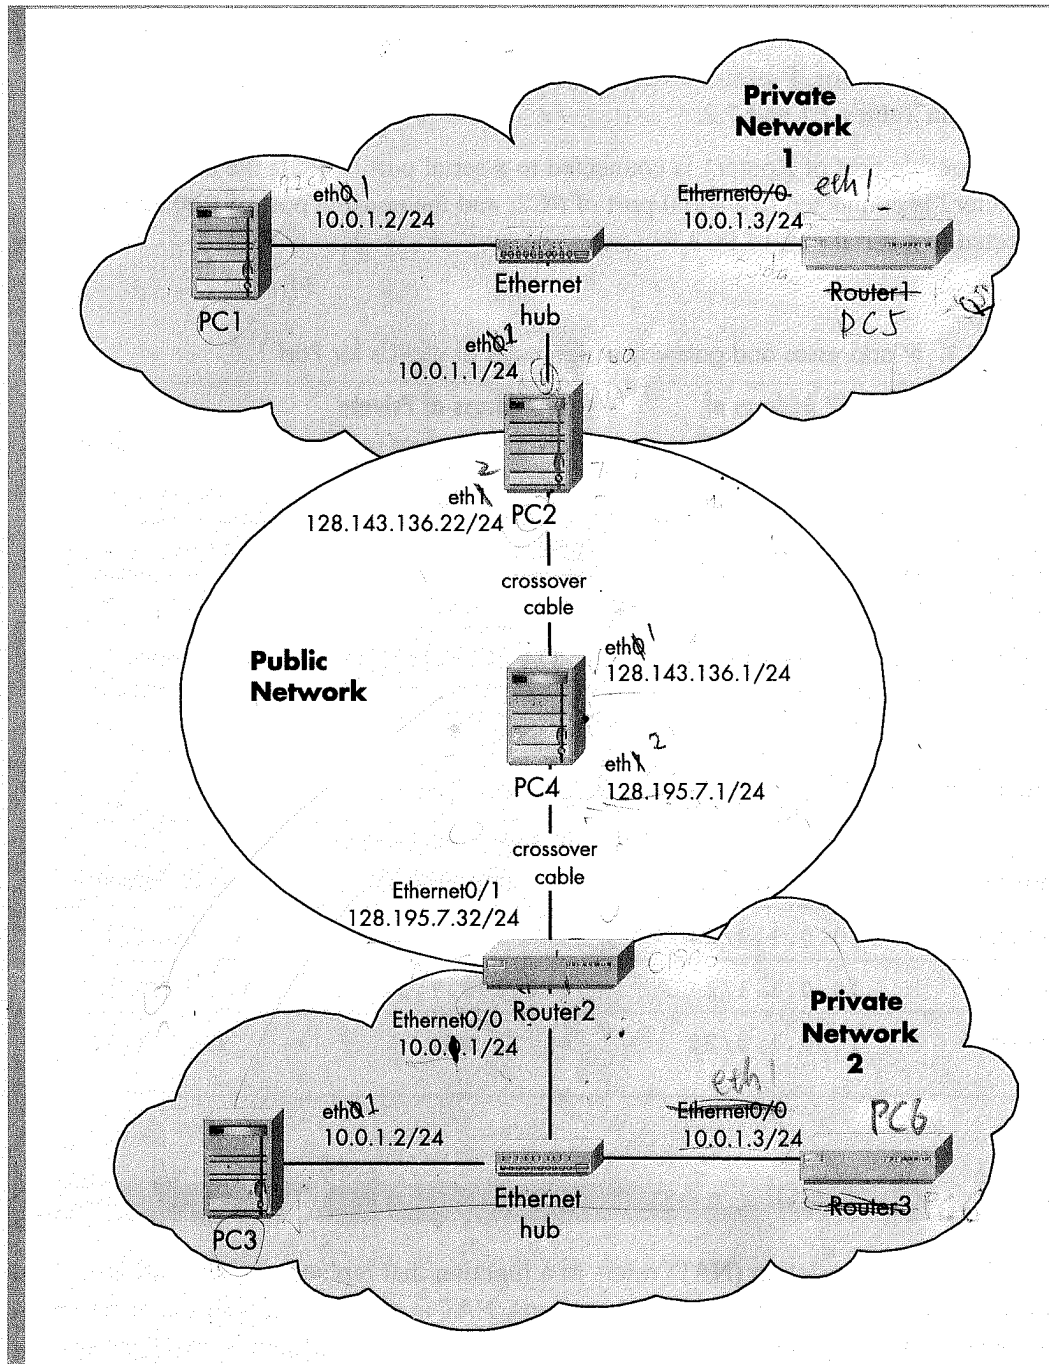


FIGURE 7.1.
Configuration for Part 1.

Note: In the private networks in Figure 7.1, Router1 and Router3 are used to mimic hosts (i.e., they are not configured to act as IP routers).

- In this lab PC2 and Router2 are routers that provide the gateways between the private and the public networks. Both PC2 and Router2 are configured as NAT devices.
- On PC2 the kernel is built with *netfilter*, an extension to the Linux kernel that provides the kernel with the ability to set IP packet filters, including NAT functions. On Router2, you will use Cisco IOS commands to configure NAT rules.
- PC4 runs as an IP router. (We use a Linux PC instead of a Cisco router so that *ethereal* can be used to capture traffic on the public network.)
- The assignment of IP addresses and default gateways for all PCs and routers are shown in Tables 7.1 and 7.2.
- The console port of Router1 is connected to a serial port of PC1; the console port of Router2 is connected to a serial port of PC2; and the console port of Router3 is connected to a serial port of PC3.

TABLE 7.1. IP addresses and gateway assignments of all PCs for Part 1.

Machine	IP Address of Interface (<i>eth0</i>) ¹	IP Address of Private Network Interface (<i>eth1</i>) ²	Default Gateway
PC1	10.0.1.2/24	none	10.0.1.1
PC2	10.0.1.1/24	128.143.136.22/24	128.143.136.1
PC3	10.0.1.2/24	none	10.0.1.1
PC4	128.143.136.1/24	128.195.7.1/24	none

TABLE 7.2. IP addresses and gateway assignments of routers for Part 1.

Machine	IP Address of Interface (<i>Ethernet0/0</i>)	IP Address of Interface (<i>Ethernet0/1</i>)	Default Gateway
PC5 Router1	10.0.1.3/24 <i>eth1</i>	none <i>eth2</i>	10.0.1.1
CISCO Router2	10.0.1.1/24	128.195.7.32/24	128.195.7.1
PC6 Router3	10.0.1.3/24 <i>eth1</i>	none <i>eth2</i>	10.0.1.1

PART I. NETWORK ADDRESS TRANSLATION (NAT)

Network address translation (NAT) refers to a function that replaces the IP addresses (and possibly the port numbers) of IP datagrams. NAT is run on routers that connect private networks to the public Internet, to replace the IP address-port pair of an IP packet with another IP address-port pair. Generally, the operations of NAT are specified in terms of a set of rules which determines how IP addresses are to be replaced.

Often, a NAT device is referred to as a *NAT box*. One of the reasons for using NAT is that it conserves IP addresses. NAT allows hosts in a private network to share public IP addresses or to limit the use of public IP addresses to a small number of hosts in the private network.

Private networks may have IP addresses that are non-Internet routable, as specified in RFC 1918. This means that regular IP routers do not have entries in their routing tables for these addresses.

In the network in Figure 7.1, both PC2 and Router2 will be configured as NAT devices. With NAT, the hosts in the private networks can access the public network (i.e., they are able to reach the addresses on the 128.143.136.0/24 and 128.195.7.0/24 networks).

EXERCISE 1(A).

Network setup.

Configure the network in Figure 7.1 with the IP address configuration shown in Tables 7.1 and 7.2. The following commands review the steps involved in the configuration.

1. On the Linux PCs, use `ifconfig` to configure the IP address of the interfaces. Add a default gateway on each PC with the command (shown for PC1):

```
PC1% route add default gw 10.0.1.1
```

2. IP forwarding must be enabled on PC2 and PC4.
3. Use a serial cable to connect a serial port of a PC to the console port of a router. Use the `kermit` command to access the routers.
4. Configure the IP addresses of interfaces Ethernet0/0 and Ethernet0/1 on the routers, and set the default gateways as shown in Table 7.2. Following is the sample configuration for Router2, where the Ethernet interface is located in slot 0:

```
Router2> enable
Password: <enable secret>
Router2# configure terminal
Router2(config)# no ip routing
Router2(config)# ip routing
Router2(config)# ip route 0.0.0.0 0.0.0.0 128.195.7.1
Router2(config)# interface Ethernet0/0
Router2(config-if)# no shutdown
Router2(config-if)# ip address 10.0.1.1 255.255.255.0
Router2(config-if)# interface Ethernet0/1 1
Router2(config-if)# no shutdown
Router2(config-if)# ip address 128.195.7.32 255.255.255.0
Router2(config-if)# end
```

In the above commands `ip route 0.0.0.0 0.0.0.0 128.195.7.1` sets the default gateway of Router2.

After completing the setup of the configuration you should be able to issue successful `ping` commands (between hosts in the private network and between hosts in the

public network). However, ping commands across a private/public network boundary are not successful.

EXERCISE 1(B).

Configuration of NAT on a Cisco router.

A Cisco router can be set up to run as a NAT device.

In Cisco IOS, the private network is referred to as *inside*, and the public network is referred to as *outside*. An IP address that is seen by hosts on the inside is called a *local address*, and an IP address that is seen by hosts on the outside is called a *global address*. There are four different types of addresses:

- An *inside local address* is an address in the private network that is not visible in the public network.
- An *inside global address* can be used in the public network for devices in the private network.
- An *outside global address* is an address in the public network that is not made known in the private network.
- An *outside local address* is an IP address that is assigned in the private network to designate a host in the public network.

Using this terminology, a NAT device translates in outgoing packets inside local addresses to outside global addresses and in incoming packets translates outside global addresses to inside local addresses.

1. **Modify the NAT table of Router2:** Use the following commands to set up Router2 as a NAT device:

- A NAT rule is added so that the private IP address of PC3, 10.0.1.2, is translated to the public address 200.0.0.2. See Table 7.3.

TABLE 7.3. Private and public address of PC3.

Machine	Inside Local Address	Outside Global Address
PC3	10.0.1.2/24	200.0.0.2/24

The IOS commands are as follows

```
Router2> enable
Password: <enable secret>
Router2# show ip nat translations
Router2# configure terminal
Router2(config)# interface Ethernet0/0
Router2(config-if)# ip nat inside
Router2(config-if)# interface Ethernet0/1
Router2(config-if)# ip nat outside
```



```
Router2(config-if)# exit
Router2(config)# ip nat inside source static 10.0.1.2
                200.0.0.2
Router2(config)# end
Router2# show ip nat translations
```

- b. After the preceding rule has been entered, display the content of the NAT table and save it to a file.



The commands used above are explained as follows.

IOS MODE: PRIVILEGED EXEC

Show ip nat translations
Displays the content of the NAT table.

IOS MODE: INTERFACE CONFIGURATION

ip nat inside
Specifies that an interface is connected to the private network.

ip nat outside
Specifies that an interface is connected to the public network.

IOS MODE: GLOBAL CONFIGURATION

ip nat inside source static *IPaddr1* *IPaddr2*
Adds a rule so that the private IP address *IPaddr1* is mapped to a public IP address *IPaddr2*.

For example, the command,
ip nat inside source static 10.0.1.2 200.0.0.2
maps the private address 10.0.1.2 to the public address 200.0.0.2.

Dynamic NAT is an alternative to the static NAT table entries used in this exercise. With dynamic NAT, a pool of global addresses is specified at the NAT device. Addresses from the pool are dynamically mapped to the private addresses whenever there is a demand for a new address.

2. **Update routing tables:** Add static routing entries to the routing table of PC4, so that traffic with destination IP address 200.0.0.0/24 is forwarded to Router2.
3. **Observe traffic at a NAT device:** To observe the IP address translation, issue ping commands between machines in the public and private network. Use *ethereal* to capture packets on the private and public interfaces of Router2.
 - a. Start an *ethereal* session on PC3 to capture the traffic from Router2 on the private network.
 - b. Start an *ethereal* session on interface *eth1* of PC4 to capture the traffic from Router2 on the public network.

c. Issue the following ping commands:

On PC3:

```
✓ PC3% ping -c 3 10.0.1.3
✓ PC3% ping -c 3 128.143.136.1
```

On Router3:

```
Router3# ping 10.0.1.2
Router3# ping 128.143.136.1
```

On PC4:

```
PC4% ping -c 3 10.0.1.2
PC4% ping -c 3 200.0.0.2
```

d. Save the *ethereal* data to files. Observe which ping commands succeed.

4. **Add additional NAT table entries:** Add NAT rules to Router2, so that Router2 and Router3 (on interface Ethernet0/0) are addressable from the public network. The private and public addresses are given in Table 7.4.

TABLE 7.4. Private and public addresses of Router2 and Router3.

Machine	Inside Local Address	Outside Global Address
Router2	10.0.1.1	200.0.0.1
Router3	10.0.1.3	200.0.0.3

Issue ping commands to verify that the IP addresses are correctly mapped.

Lab Report

- Include the NAT table of Router2 and provide an explanation of the columns of the table.
- For each of the preceding ping commands, provide an explanation of why the command succeeds or fails.
- Include the IP source address and IP destination address from the IP header data of an ICMP Echo Request and the corresponding ICMP Echo Reply packet before and after it passes through Router2.

EXERCISE 1(B).

IP masquerading with a Linux PC.

In this exercise we consider a special use of NAT that allows multiple private IP addresses to be mapped to a single public IP address. This use of NAT is called *IP masquerading*, *port address translation (PAT)*, or *network address and port translation (NAPT)*. Here, the private network has only a single public IP address but has multiple hosts in the private network. IP masquerading modifies the port number of packets so that the single public IP address can be overloaded.

In this exercise, PC2 will be configured to perform IP masquerading. The Linux kernel on all PCs has been built with *netfilter*, which adds the ability to set IP packet filters in a Linux system. IP packet filters are used to add firewalls as well as NAT functionality to a system. The `iptables` command is used to set up, maintain, and inspect IP packet filter rules to a Linux kernel.

On a Linux system, the configuration of NAT manipulates a set of rules of the *netfilter* utility, called the NAT table. The rules in the NAT table are grouped in so-called chains. Two of the built-in chains are called `PREROUTING` and `POSTROUTING`:

- **PREROUTING:** The rules in this chain are applied to incoming datagrams.
- **POSTROUTING:** The rules in this chain are applied to outgoing datagrams. The main rule is `SNAT` (source network address translation), which specifies how the source address of an outgoing IP datagram should be modified.

Commands that manipulate the NAT table start with

```
PC2% iptables -t nat
```

The following are some of the most important commands that manipulate the NAT table.

```
iptables -t nat -L
```

Displays all rules in the NAT table.

```
iptables -t nat -D POSTROUTING 1
```

Deletes the first rule in the `POSTROUTING` chain of the NAT table

```
iptables -t nat -F
```

Deletes all entries in (flushes) the NAT table

```
iptables -t nat -A POSTROUTING -j SNAT --to IPAddr1 -s IPAddr2/netmask
```

Adds the following rule to the `POSTROUTING` chain of the NAT table:

"in IP datagrams that go to the public network, the IP source address `IPAddr2/netmask` is changed to `IPAddr1`".

For example, in the command

```
iptables -t nat -A POSTROUTING -j SNAT --to 128.195.7.32 -s 10.0.1.0/24
```

the source address of outgoing IP datagrams that match `10.0.1.0/24` is changed to `128.195.7.32`.

1. **Modify the NAT table of PC2:** On PC2, add a rule to the NAT table so that the IP source address of all outgoing IP datagrams are set to IP address `128.143.136.22`.

Display the content of the NAT table and save it to a file.

2. **Observe traffic at a NAT device:**

- a. To observe the IP address translation, observe traffic on both interfaces of PC2 that are between the private networks and the public network. On PC2, run *ethereal* on both `eth0` and `eth1`, and start to capture traffic.



- b. Establish a set of *Telnet* sessions and log in to remote machines, using the following *telnet* commands:

On PC1:

```
✓PC1% telnet 10.0.1.3
✓PC1% telnet 128.143.136.1
```

On Router1:

```
✓Router1# telnet 10.0.1.2
✓Router1# telnet 128.143.136.1
```

On PC4:

```
✓PC4% telnet 10.0.1.2
```

When opening a *Telnet* session to a Cisco router, IOS asks for a login password. This password is generally not the same password as the enable secret, which is used to enter the privileged EXEC mode. However, in the Internet lab, the login password for *Telnet* sessions is set to use the same password as the enable secret.

- c. Save the *ethereal* data to files. Observe which *telnet* commands succeed.
- d. For the successful *Telnet* sessions, observe how the IP addresses and port numbers are mapped.

3. **Observe mapping of ICMP packets:** The *ping* command sends out ICMP Echo Request messages and receives ICMP Echo Reply messages. Since ICMP messages do not contain a port number, it is not entirely obvious how a NAT device that performs IP masquerading can direct ICMP Echo Reply messages that return from the public network to the private network. In this exercise, you will explore how a NAT device handles ICMP messages.

- a. On PC2, run *ethereal* on both *eth0* and *eth1*. Use the appropriate filters to capture the traffic generated by *ping* commands.
- b. Issue the following *ping* commands:

On PC1:

```
✓PC1% ping -c 3 10.0.1.3
✓PC1% ping -c 3 128.143.136.1
```

On Router1:

```
✓Router1# ping 10.0.1.2
✓Router1# ping 128.143.136.1
```

On PC4:

```
✓PC4% ping -c 3 10.0.1.2
```

- c. Save the *ethereal* output and the output of *ping* commands into files.

Lab Report

- For each of the preceding *telnet* and *ping* commands, provide an explanation why a command succeeds or fails.

- For each successful *Telnet* session, include the IP header data of an outgoing and an incoming packet header (with respect to the private network).
- For each successful *ping* command, include the IP header data of an outgoing ICMP Echo Request message and an incoming ICMP Echo Reply message (with respect to the private network).
- How does a PC know that a packet coming from the public network is destined to a host in the private network?
- Explain the steps performed by the Linux kernel during IP address translation.

EXERCISE 1(D).**NAT and FTP.**

NAT can create problems for applications, which carry the IP addresses in the payload of an IP datagram. An example of such an application is the File Transfer Protocol (*FTP*).

In this exercise, you establish an *FTP* session from PC3 in the private network to PC2 in the public network and observe how *FTP* works with NAT.

1. Start *ethereal* on interface *eth0* of PC4 and on interface *eth0* of PC3.
2. **FTP session between two hosts in the public network:**
 - a. Start an *FTP* session from PC4 to PC2 (the `-d` option prints out debug messages):


```
PC4% cd /root/labdata
PC4% ftp -d 128.143.136.22
```

Log in with user name *root* and enter the root password.
 - b. Download a file from the *FTP* server:


```
ftp> get fname
```

where *fname* is a file on the remote server. (You can use the command `ls` to obtain a list of all files in the remote directory.)
 - c. Use the traffic captured by *ethereal* to determine where the *payload* of *FTP* data carries information on IP addresses.
 - d. Save the *ethereal* output and the *FTP* debug information output into files.
3. **FTP session from a private to the public network:**
 - a. Use the same commands as previously to download a file from PC2 to PC3:


```
PC3% ftp -d 128.143.136.22
```

Is the *FTP* session establishment successful?
 - b. Save the traffic captured by *ethereal* and save the *FTP* debug information output. Make sure that you save enough data to answer the lab report questions.

Lab Report Use the captured data to explain the outcome of the *FTP* experiment. In particular, if the file was successfully *downloaded*, explain how the problem of sending the IP address as part of the data payload of the IP packet is solved.



You need to do this part

PART 2. DYNAMIC HOST CONFIGURATION PROTOCOL (DHCP)

The Dynamic Host Configuration Protocol (DHCP) can be used to dynamically set and change configuration parameters of Internet hosts, including IP address, netmask, default router, and DNS server. DHCP is based on a client-server model. DHCP clients send requests to a DHCP server and the server responds with an allocation of IP addresses and other configuration parameters.

In this part of the lab, you will also learn about DHCP relay agents. When the DHCP client and DHCP server are not on the same IP network, DHCP relay agents can act as routers of DHCP messages. A DHCP relay agent can forward DHCP requests from a DHCP client to a DHCP server and it can forward the reply messages from the DHCP server to the DHCP client.

The network configuration for Part 2 is shown in Figure 7.2. PC1, PC3, and PC4 are set up as DHCP clients and initially do not have IP addresses. PC2 is configured as a DHCP server, which listens for DHCP requests on all of its interfaces and transmits network configuration parameters. Router1 acts as a DHCP relay agent, which forwards DHCP messages between different IP networks.

Table 7.5 lists the range of addresses that are associated at the DHCP server PC2 with each IP network.




FIGURE 7.2.
Network configuration for Part 2.

TABLE 7.5. Configuration of the PCs in Part 2.

Machine	IP Address of Internet Interface (<i>eth0</i>) ¹	IP Address of Private Network Interface (<i>eth1</i>) ²	Default Gateway
PC1	none	none	none
PC2	10.0.1.21/24	10.0.2.21/24	10.0.2.1
PC3	none	none	none
PC4	none	none	none

EXERCISE 2(A).
Network setup.

1. We strongly recommend that you reboot the PCs and the routers before you proceed.
2. Set up the network topology as shown in Figure 7.2. Configure the IP addresses of the PCs and Router1 as shown in Table 7.6 and 7.7.

TABLE 7.6. Router configuration

Machine	IP Address of Interface (<i>Ethernet0/0</i>)	IP Address of Interface (<i>Ethernet0/1</i>)	Default Gateway
Router1	10.0.2.1/24	10.0.3.1/24	10.0.2.21

TABLE 7.7. DHCP server configuration.

Subnet	Range of Addresses	Default Router
10.0.1.0/24	10.0.1.2 to 10.0.1.10	10.0.1.21
10.0.3.0/24	10.0.3.2 to 10.0.3.10	10.0.3.1

3. It is important that PC1, PC3, and PC4 do not have a default route and do not have an IP address associated with their respective interface *eth0*.

Review the routing table and the interface configuration. On PC1, this is done with the commands

```
PC1% netstat -rn
PC1% ifconfig -a
```


In Linux, routing tables display the default route as an entry with destination 0.0.0.0. If the routing table shows a default route, you can delete this and all other routing table entries by setting the IP address to 0.0.0.0. This is done with the following command:

```
PC1% ifconfig eth0 0.0.0.0 up
```

the path should be
/etc/dhcp/dhcpd.conf

EXERCISE 2(B).

Configuring and starting a DHCP server.

On a Linux system, a DHCP server is started with the command `dhcpd`. The DHCP server reads the configuration file `/etc/dhcpd.conf`. The configuration file contains information on available IP addresses, as well as other configuration information. The following is an example of a configuration file for a DHCP server:

```
#dhcpd.conf file
ddns-update-style none;
default-lease-time 600;
subnet 10.0.1.0 netmask 255.255.255.0 {
    range 10.0.1.10 10.0.1.100;
    option routers 10.0.1.1;
    default-lease-time 120;
}
subnet 10.0.2.0 netmask 255.255.255.0 {
    range 10.0.2.101 10.0.2.200;
}
subnet 10.0.3.0 netmask 255.255.255.0 {
    range 10.0.3.6 10.0.3.10;
}
```

Hint:

To bind interface to a subnet, use interface command. eg:

```
subnet 10.0.2.0 netmask 255.255.255.0 {
    range 10.0.2.101 10.0.2.200;
    interface p1p2;
}
```

This binds p1p2 to subnet 10.0.2.0

The DHCP client is assigned an IP address for a period of time that is known as a lease. The preceding configuration file assigns IP addresses for a lease time of 600 seconds (default-lease-time). For requests on network 10.0.1.0/24, the DHCP server assigns IP addresses in the range 10.0.1.10–10.0.1.100, assigns 10.0.1.1 as the default gateway, and limits the lease of addresses to 120 seconds, thus overruling the global limit of 600 seconds. For requests on network 10.0.2.0/24, the server assigns IP addresses in the range 10.0.2.101–10.0.2.200.

- 1. Set the DHCP configuration file:** On PC2 set up the configuration file so that IP addresses are assigned as follows: On network 10.0.1.0/24, IP addresses are assigned in the range 10.0.1.2–10.0.1.10 with default gateway 10.0.1.21. On network 10.0.3.0/24, IP addresses are assigned in the range 10.0.3.2–10.0.3.10 with default gateway 10.0.3.1.

Note that these assignments are similar to, but not identical with the configuration file shown above.

- 2. Start the DHCP server:** On PC2, start the DHCP server by typing

```
PC2% dhcpd
```

The DHCP server listens for requests from DHCP clients on all its interfaces. In Linux, the DHCP server must be restarted each time the configuration file is modified. Since only one DHCP server can run at a time, you may need to terminate the current DHCP server process.

EXERCISE 2(C).

Starting a DHCP client.

The following steps start a DHCP client on PC1:

1. Preparation: On PC1, perform the following functions:

- a. Ensure that no default router entry exists in the routing table.
- b. A Linux DHCP client caches information from previous uses of DHCP. The cached information is stored in files `dhclient`

```
/var/lib/dhcp/dhclient.leases
/var/lib/dhcp/dhclient.leases~
```

Since this cached information may interfere with your work, delete these files, if they exist.

- c. Start *etherreal* on interface *eth0* of PC2. (Set the display filter to `bootp.dhcp` so that only DHCP traffic is displayed in the window.)

2. Start a DHCP client: The DHCP client on PC1 is started with the command

```
PC1% dhclient -d eth0
```

Save the data that is captured by *etherreal* to a file. Save enough data to answer the following questions from the captured traffic:



- a. Which IP address is assigned to PC1?
 - b. Observe the source and destination IP addresses of the packets that are sent between DHCP client and DHCP server.
 - c. How is it possible that a host can send and receive DHCP packets, even though it does not have an IP address?
 - d. Do you observe any ARP packets? If so, explain the function of ARP in this context.
 - e. Observe and interpret the output of the DHCP packets. You should see the following packet types: DHCP Discover, DHCP Offer, DHCP Request, DHCP ACK.
 - f. Identify and interpret all option fields in the DHCP packet types that you observe.
- #### 3. Renewing leases of IP addresses: The DHCP client is assigned an IP address for a limited period of time, which is called a *lease*. The maximum time of a lease is specified in the file *dhcpd.conf*. Information on current leases is stored at both the client side and the server side.

- a. In Linux, information on the current leases is stored in the following files:

At the DHCP server: `/var/lib/dhcp/dhcpd.leases`

At the DHCP client: `/var/lib/dhcp/dhclient.leases`

dhcpd

dhclient

- b. To interpret the content of the files, refer to the manual pages of `dhcpd.conf`, `dhcp-options`, and `dhcpd.leases`.

- c. Save the files that contain the information on current leases.

- d. Observe how a DHCP client renews a lease and saves the captured traffic to a file.

- What type of DHCP message can be observed?
- How long does a DHCP client wait until it attempts to renew its lease?

- e. Stop the process that runs the DHCP server by terminating the process `dhcpd` with the command

```
PC2% pkill dhcpd
```

Observe what the DHCP client does when it cannot reach the DHCP server. Use the command `ifconfig -a` to see how long the DHCP client waits until it releases the leased IP address.

- f. Restart the DHCP server process by typing

```
PC2% dhcpd
```

- 4. Starting more DHCP clients:** Repeat the instructions in Step 2 and start DHCP clients on PC3 and PC4.

- a. The expected outcome is that PC4 receives an IP address but that PC3 is not successful. Why is the negative outcome for PC3 expected?
- b. Compare the IP addresses assigned to PC1 and PC4. Is there a specific order in which IP addresses are assigned by the DHCP server?

Lab Report Answer the following questions using the *ethereal* output:

- Use a figure to explain the packets that were exchanged by the DHCP client and the DHCP server as part of the process of acquiring an IP address.
- Explain the entries in the lease file `dhcpd.leases`. How is the content of the lease file used when a DHCP client cannot contact the DHCP server?
- In most client-server applications, the port number of a server is a well-known number (e.g., an *FTP* server uses port number 21, the *Telnet* server uses port number 23, etc.), while the client uses a currently available (ephemeral) port number. DHCP is different. Here, both the client and the server use a well-known port: UDP port 67 for the DHCP server and UDP port 68 for the DHCP client. Refer to RFC 2131 and provide an explanation for this protocol design choice.

EXERCISE 2(D).**DHCP relay agent.**

A DHCP relay agent can forward DHCP packets when the DHCP server and the DHCP client are not on the same network. Note that the role of a DHCP relay agent is not entirely trivial, since it acts as a router for a host that does not have an IP address. Here you explore how packets from the client reach the server on another network and how the response from the server reaches the DHCP client.

The DHCP server is configured to allocate addresses as shown in Table 7.7.

Next we set up a Cisco router as a DHCP relay agent.

- 1. Setting up a Cisco router as a DHCP relay agent:** The following commands set up Router1 as a DHCP relay agent. In essence, Router1 is configured to forward UDP packets.

Start the DHCP relay agent on Router1 as follows:

```
Router> enable
Password: <enable secret>
Router1# configure terminal
Router1(config)# ip forward-protocol udp
Router1(config)# interface Ethernet0/1
Router1(config-if)# ip helper-address 10.0.2.21
Router1(config-if)# end
```

IOS MODE: GLOBAL CONFIGURATION

```
ip forward-protocol udp
  Enables UDP packet forwarding.
```

IOS MODE: INTERFACE CONFIGURATION

```
ip helper-address Ipaddr
  Forwards DHCP request packets received on the current interface to IP address
  IPaddr, which should be the IP address of a DHCP server.
```

- 2.** Start *etherreal* on PC2 and PC3.
- 3.** Make sure that the DHCP server is running on PC2. If necessary, start a new DHCP server.
- 4.** Start a DHCP client on PC3 with
PC3% `dhclient -d eth0`
- 5.** Verify that an IP address has been assigned to PC3.

According to the configuration file, the DHCP configuration on network 10.0.2.0/24 does not set a default router. Verify that this is correct by inspecting the routing table.

6. Answer the following questions:

- Does the DHCP relay agent modify DHCP packets or the IP header? If so, what are the modifications?
- How does the relay agent redirect the replies from the DHCP server? Does it broadcast them or unicast them to the DHCP client?
- Is there a difference in the response of the DHCP server as compared to the DHCP configuration of PC1? If so, explain the difference.
- How does the DHCP server (PC2) know on which network PC3 is located when it receives the DHCP request?
- What is the destination IP address of the first DHCP packet that the DHCP server sends to PC3?

7. Save the DHCP packets that are captured by *ethereal* to a file.**Lab Report**

- Include the *ethereal* data of the first three DHCP packets that are exchanged between PC3 and PC2.
- Provide answers to the questions in Exercise 2(D).
- What happens if a network has multiple DHCP servers?

Again, this part is not required

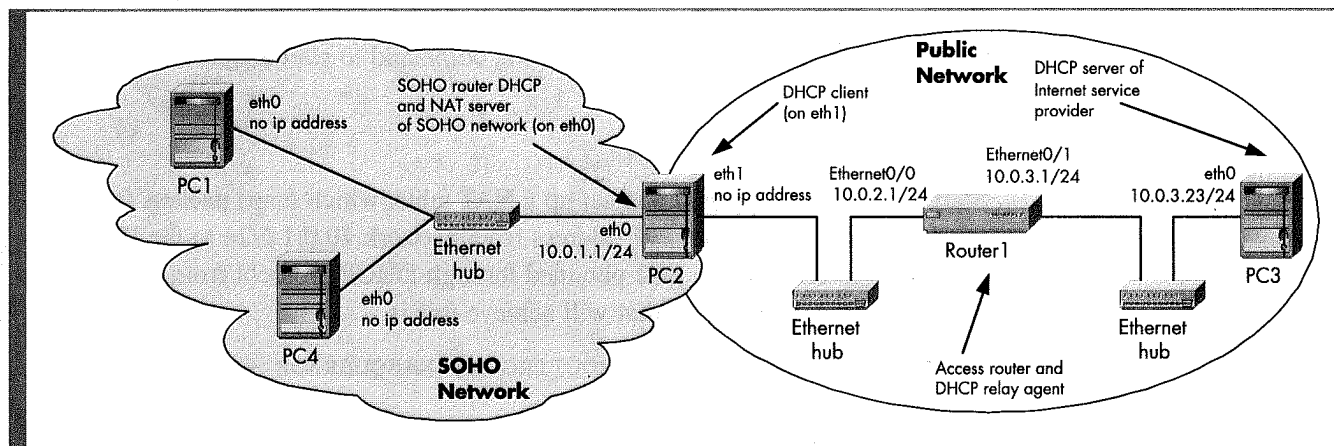
PART 3. COMBINING NAT AND DHCP

Figure 7.3 shows a network configuration that can be found in many SOHO (small office, home office) networks.

- The SOHO network is a private network with multiple hosts (PC1 and PC4) and one IP router (PC2).
- The IP router of the SOHO network (SOHO router) provides access to the public Internet by connecting to a router of an Internet service provider. The SOHO router obtains a single IP address on the “public” interface of the SOHO network via DHCP from a DHCP server (PC3) of the Internet service provider.
- The SOHO router (PC2) works as a DHCP server and NAT server for the hosts in the SOHO network.

In this network setup, all SOHO hosts can share a single public IP address, which is dynamically assigned by the Internet service provider. Furthermore, the SOHO network requires minimal IP configuration. The hosts in the SOHO network obtain their IP addresses from the SOHO router. The SOHO router obtains its (public) IP address from the Internet service provider.

Your task is to set up the entire SOHO network, including the router and the DHCP server of the Internet service provider.

**FIGURE 7.3.**

Network configuration for Part 3.

EXERCISE 3:

The network configuration is shown in Figure 7.3. (The connections of the cables are identical to Figure 7.2). To reset the configuration of all machines, we recommend rebooting the PCs and the router.

1. DHCP server: PC3 is the DHCP server of the Internet service provider.
 - a. Configure PC3 with IP address 10.0.3.23/24 on interface *eth0* and with default gateway 10.0.3.1.
 - b. Configure and start a DHCP server on PC3. On PC3, set up the configuration file so that IP addresses in the range 10.0.2.2–10.0.2.10 are assigned for requests on network 10.0.2.0/24, and addresses in the range 10.0.3.2–10.0.3.10 are assigned for requests on network 10.0.3.0/24.
2. Router and DHCP relay agent: Router1 is the IP router to which the SOHO network sends its external traffic. Also, Router1 is a DHCP relay agent.
 - a. Configure Router1 with IP addresses 10.0.2.1/24 on interface *Ethernet0/0* and 10.0.3.1/24 on interface *Ethernet0/1*.
 - b. The routing table of Router1 should reflect that all traffic to network 10.0.2.0/24 is sent on interface *Ethernet0/0*, and all other traffic is sent on interface *Ethernet0/1*.
 - c. Configure Router1 as a DHCP relay agent, so that requests from DHCP client PC2 reach DHCP server PC3.
3. SOHO router: PC2 is the SOHO router.
 - a. Set up PC2 so that it is a DHCP client on interface *eth1*.
 - b. Set up PC2 as an IP router. That is, IP forwarding must be enabled. The routing table entries must reflect that traffic to network 10.0.1.0/24 must be routed on interface *eth0*, and all other traffic must be sent to Router1 at 10.0.2.1.

- c. Configure PC2 as DHCP server on interface *eth0* for addresses in the range 10.0.1.2–10.0.1.10. Execute the following command to start a DHCP server process on PC2:

```
PC2% dhcpd eth0
```

- d. Start a NAT server on PC2 and set up a NAT table, which maps packets from the SOHO network with source IP address from network 10.0.1.0/24 to the IP address of interface *eth1*, PC2 obtained through DHCP protocol from PC3. The command for adding a rule that will achieve this is

```
iptables -t nat -A POSTROUTING -j MASQUERADE -o eth1 -s
10.0.1.0/24
```

4. Hosts in PCs: PC1 and PC4 are hosts in the SOHO network. Set up PC1 and PC4 as DHCP clients on interfaces *eth0*.

5. Collecting the results:



- a. Display the routing tables from all PCs with `netstat -rn` and the IP configuration with `ifconfig -a` and save the results.

What are the IP addresses assigned to PC1 and PC4? How are the IP addresses mapped to the public IP address defined on the NAT server PC2?



- b. Display and save the NAT table of PC2.
 c. Start *ethereal* on PC1 (*eth0*), PC2 (*eth1*), and PC3 (*eth0*).
 d. Issue a `ping` command from PC1 to PC3:

```
PC1% ping -c 5 10.0.3.23
```



- e. Save the traffic captured by *ethereal* on one of the PCs to a file.



Lab Report

- Include the *ethereal* data from the first ICMP Request and ICMP Reply messages.
- Include the routing table and the output of the `ifconfig` command from all PCs.
- Include the NAT table from PC2.