CHAPTER 5

CONFIGURING AND MANAGING THE NETWORK INFRASTRUCTURE

Chapter Outline

Objectives

- Describe the steps for obtaining a domain name
- Describe the steps for getting IP addresses assigned to your network
- Understand the purpose of the DHCP server
- Describe the use of NAT and PAT and how these technologies differ
- Understand the function of the DNS server
- Describe the purpose of SNMP

Key Terms

IANA
gTLDs
ccTLDs
.int
in-addr.arpa
RIRs
AS
ICANN
ARIN
domain registrars
whois protocol
BOOTP
DHCP
Lease Time
DHCP Discover
DHCP Offer
DHCP Request
DHCP ACK
ipconfig/release
ipconfig /renew
Automatic Private IP addressing (APIPA)
Unicast
ip helper [*ip address of the DHCP server*]
MT Discover
MT Offer
MT Request
MT ACK
ARP broadcast
SOHO
binding
Network Address Translation (NAT)
Port Address Translation (PAT)
local address
global address
static NAT
dynamic NAT
NAT overload
DNS
Forward Domain Name Service
Reverse Domain Name Service
Root Hints file (root.hints)
TLD
Country Domain
Dig (Domain Information Groper)
nslookup
authoritative name server
non-authoritative answer
FQDN
PQDN
RR
SOA
A record (Address record)
PTR record (Pointer record)
CNAME record (Canonical name record)
NS record
MX record
TXT record
SPF
SRV record

INTRODUCTION

This chapter examines the issue of configuring and managing the network infrastructure. Section 5-1 addresses the procedure used to obtain a domain name for your network and the steps required to get IP addresses assigned to your network. After you have an IP address assigned to your network, you will have to manage the IP assignment.

Section 5-2 discusses IP management with DHCP. DHCP simplifies the steps for IP assignment by dynamically assigning IP addresses as they are needed. This section also examines the deployment of the DHCP relay function that is provided by a router.

Public IP addresses are a commodity and the network administrator must be aware of the techniques used to properly manage the existing pool of IP addresses. This is discussed in Section 5-3, which examines the techniques for scaling the network with NAT and PAT. The steps for configuring NAT and PAT on a router are also presented.

Section 5-4 looks at the Domain Name Service (DNS). The purpose of the DNS server is to translate a human readable name to an IP address or an IP address to a domain name. This section also examines the DNS tree hierarchy and the many DNS services currently available.

IANA
Internet Assigned Numbers Authority, which is responsible for the global coordination of the DNS Root, IP addressing, and other Internet Protocol resources.

gTLDs
Generic (g) top-level domains. Includes .com, .net, .org, and .info.

ccTLDs
Country-code (cc) top-level domains. Includes .us, .uk, .ca, and .au.

.int
Intergovernmental domain registries is used for registering organizations established by international treaties between or among national governments.

in-addr.arpa
The reverse DNS lookup for IPv4 addresses on the Internet.

5-1 DOMAIN NAME AND IP ASSIGNMENT

There are two key elements used by the general population when accessing websites on the Internet. One is the Internet name of the website and the other is its public IP address. These two elements go hand in hand. People generally connect to Internet services via Internet hostnames, (for example, www.example.com), but behind the scenes, the Internet name is translated to a public IP address. Both the IP address assignment and the Internet domain name are governed at the highest level by the Internet Assigned Numbers Authority (**IANA**).

IANA is one of the Internet's oldest organizations and was set up to be in charge of the Internet management authorities or registration authorities. IANA has three primary functions:

1. **Domain name management**: IANA manages the DNS root zone for the generic (g) top-level domains (**gTLDs**), such as .COM, .NET, .ORG, .INFO, and country-code (cc) top-level domains (**ccTLDs**), such as .US, .UK, and .AU. IANA maintains the **.int** (intergovernmental) domain registries, which are exclusive registrations for intergovernmental treaty organizations, such as the United Nations (un.int) and NATO (nato.int), Asnthe.int, .arpa domains, and an IDN practices resource. IANA maintains the .arpa domain registries, which include the in-addr.arpa domain. The **in-addr.arpa** is the reverse DNS lookup for IPv4 addresses on the Internet. IANA also maintains the IDN (Internationalized

Domain Name) practices repository known as the language table registry. This allows for domain name registration containing international characters (for example, müller.info).

2. **Number resources management**: IANA coordinates the global pool of IP addresses, which include both IPv4 and IPv6. To coordinate the global effort of IP address allocation more effectively, IANA delegates the allocation to the regional Internet registries (**RIR**), each of which is responsible for a different area. The five RIRs accounting for the different regions of the world are as follows:
   - **AfriNIC**: Africa Region
   - **APNIC**: Asia/Pacific Region
   - **ARIN**: North America Region
   - **LACNIC**: Latin America and some Caribbean Islands
   - **RIPE NCC**: Europe, the Middle East, and Central Asia

   IANA is also responsible for the **AS** (Autonomous System) number allocation, which is used in BGP to route Internet traffic. This allocation is delegated to the RIRs the same as the IP address allocation.
3. **Protocol Assignments**: IANA is also responsible for maintaining the registries of protocol names and numbers used in the Internet today. These protocol-numbering systems are managed by IANA in conjunction with standards bodies.

RIRs
Regional Internet registries. The RIRs are the organizations that actually allocate IP addresses to ISPs.

AS
Autonomous System. These numbers are used by various routing protocols and are a collection of connected Internet Protocol (IP) routing prefixes.

Today, IANA is working under the direct support from the Internet Corporation of Assigned Names and Numbers (**ICANN**); however, these organizations do not directly allocate IP address space nor do they register domain names for the general public. In North America, IP addresses are assigned by the American Registry for Internet Numbers (**ARIN**). The web address for ARIN is http://www.arin.net. ARIN assigns IP address space to Internet service providers (ISP) and end users, but only to those who qualify. To qualify requires that the ISP or end user be large enough to merit a block of addresses.

When ARIN allocates a block of addresses to the ISP, the ISP will then issue addresses to their customers. For example, Telco could be the ISP that has a large block of IP addresses and issues an IP address to a user. It is also possible for a local ISP to be assigned a block of IP addresses from ARIN, but the local ISP must have a large number of users.

ARIN also assigns end users IP addresses. Once again, the end user must qualify to receive a block of addresses from ARIN. This usually means that the end user must be large. For example, many universities and large businesses can receive a block of IP addresses from ARIN; however, most end users will get their IP addresses from an ISP (for example, Telco) or have IP addresses assigned dynamically when they connect to the ISP.

Today, available IPv4 address space is limited. It is predicted that the available IPv4 space will be totally depleted within a few years. As a result, it has become increasingly difficult to acquire IPv4 space from ARIN. In fact, it is next to impossible to acquire Class B IP space nowadays; however, it is possible to buy a pool of IP addresses from the ISP, but the larger the IP range, the more expensive it is. There

ICANN
Internet Corporation of Assigned Names and Numbers. This organization coordinates the Domain Name System (DNS), Internet Protocol (IP) addresses, space allocation, protocol identifier assignment, generic (gTLD), country code (ccTLD), top-level domain name system management, and root server system management functions.

ARIN
American Registry for Internet Numbers. Allocates Internet Protocol resources; develops consensus-based policies; and facilitates the advancement of the Internet through information and educational outreach.

domain registrars
Has control over the granting of domains within certain top-level domains.

is good news, however. There is a big push by the Internet community to transition to IPv6. There are more abundant resources of IPv6 addresses available, and it is much easier to acquire IPv6 address space. This is being used to encourage people to make the transition to IPv6 (which is covered in Chapter 8, "IPv6").

The Internet hostname is a subset of the Internet domain name that people can identify with. For example, **www.example.com** is a web server for the domain example.com. The Internet domain name is the identity of the organization. The first step to obtain an Internet domain name is to find a domain name registrar. The **domain registrar** has control over the granting of domains within certain Top Level Domains (TLD). IANA and ICANN do not directly register domain names for the general public. ICANN delegates the top-level domain (TLD) registry to other companies or organizations. A couple of the most notable TLD registrars are Verisign, which is a company authorized to operate the TLD for .com and .net, and Educause, which is an organization operating the TLD for the .edu domain. The company Verisign, delegates the responsibilities further to other domain registrars like networksolutions.com, godaddy.com, tucows.com, and so on.

An Internet domain can be purchased from any of these registrars. When you get on the registrar's website, you will be able to input a domain name. The registrar will check whether the domain name is available. If the domain name is available, you will be prompted to complete the application for the domain name and enter the DNS servers that are to be used to host the domain. The DNS servers will be assigned an IP address and names. When the network's DNS servers are placed online, the root servers will point to the network's DNS servers. These DNS servers then become the authoritative DNS servers for the domain.

whois protocol
Queries databases that store user registration information of an Internet domain name and an IP space.

The registration for both the IP address and the Internet domain name can be verified using the **whois protocol** in a Linux environment. The *whois* protocol queries databases that store user registration information of an Internet domain name and IP space. The *whois* information gives the ownership information that includes the point of contact of a resource. There are many *whois* servers that are accessible via the web interface. All of these derive from the simple UNIX line command **whois**, which is still available today. The following example shows the result of the **whois** command entered at a UNIX prompt for the domain "example.com."

```
 [admin@noc ~]$ whois  example.com
[Querying whois.verisign-grs.com]
[whois.verisign-grs.com]

Whois Server Version 2.0

Domain names in the .com and .net domains can now be registered with
many different competing registrars. Go to http://www.internic.net for
detailed information.

   Domain Name: EXAMPLE.COM
   Registrar: RESERVED-INTERNET ASSIGNED NUMBERS AUTHORITY
   Whois Server: whois.iana.org
   Referral URL: http://res-dom.iana.org
   Name Server: A.IANA-SERVERS.NET
   Name Server: B.IANA-SERVERS.NET
```

```
   Status: clientDeleteProhibited
   Status: clientTransferProhibited
   Status: clientUpdateProhibited
   Updated Date: 14-aug-2011
   Creation Date: 14-aug-1995
   Expiration Date: 13-aug-2012

>>> Last update of whois database: Tue, 07 Feb 2012 03:55:40 UTC <<<
```

It is surprising that the domain example.com is actually a reserved domain registered by IANA. The domain has two authoritative domain name servers of a.iana-servers.net and b.iana-servers.net. The following example shows the whois result when querying an IP space of 10.0.0.0:

```
[admin@noc ~]$ whois 10.0.0.0
[Querying whois.arin.net]
[whois.arin.net]
#
# Query terms are ambiguous. The query is assumed to be:
#     "n 10.0.0.0"
#
# Use "?" to get help.
#

#
# The following results may also be obtained via:
# http://whois.arin.net/rest/nets;q=10.0.0.0?showDetails=true&showARIN
=false&ext=netref2
#

NetRange:       10.0.0.0 - 10.255.255.255
CIDR:           10.0.0.0/8
OriginAS:
NetName:        PRIVATE-ADDRESS-ABLK-RFC1918-IANA-RESERVED
NetHandle:      NET-10-0-0-0-1
Parent:
NetType:        IANA Special Use
Comment:        This block is used as private address space.
Comment:        Traffic from these addresses does not come from IANA.
Comment:        IANA has simply reserved these numbers in its database
Comment:        and does not use or operate them. We are not the source
Comment:        of activity you may see on logs or in e-mail records.
Comment:        Please refer to  http://www.iana.org/abuse/
Comment:
Comment:        Addresses from this block can be used by
Comment:        anyone without any need to coordinate with
Comment:        IANA or an Internet registry. Addresses from
Comment:        this block are used in multiple, separately
Comment:        operated networks.
Comment:
```

```
Comment:        This block was assigned by the IETF in the
Comment:        Best Current Practice document, RFC 1918
Comment:        which can be found at:
Comment:
Comment:        http://www.rfc-editor.org/rfc/rfc1918.txt
RegDate:
Updated:        2011-04-12
Ref:            http://whois.arin.net/rest/net/NET-10-0-0-0-1
OrgName:        Internet Assigned Numbers Authority
OrgId:          IANA
Address:        4676 Admiralty Way, Suite 330
City:           Marina del Rey
StateProv:      CA
PostalCode:     90292-6695
Country:        US
RegDate:
Updated:        2004-02-24
Ref:            http://whois.arin.net/rest/org/IANA

OrgAbuseHandle: IANA-IP-ARIN
OrgAbuseName:   Internet Corporation for Assigned Names and Number
OrgAbusePhone:  +1-310-301-5820
OrgAbuseEmail:  abuse@iana.org
OrgAbuseRef:    http://whois.arin.net/rest/poc/IANA-IP-ARIN

OrgTechHandle: IANA-IP-ARIN
OrgTechName:   Internet Corporation for Assigned Names and Number
OrgTechPhone:  +1-310-301-5820
OrgTechEmail:  abuse@iana.org
OrgTechRef:    http://whois.arin.net/rest/poc/IANA-IP-ARIN

#
# ARIN WHOIS data and services are subject to the Terms of Use
# available at: https://www.arin.net/whois_tou.html
#
```

The IP space of 10.0.0.0, which is a well-known private IP address range, is registered as a reserved private IP address block by IANA. The information shows the ownership information (IANA), as well as the point of contacts for the IP space. (Note: This is listed under OrgId: and Address:.) Note that the **whois** results for example.com show that the Internet domain name has an expiration date, while the IP address space 10.0.0.0 does not. The expiration date for example.com is listed as **Expiration Date: 13-aug-2012**.

Typically, an Internet domain name is purchased through a domain registrar for one year or multiple years and that sets the expiration date of the domain from the date of purchase. The domain owner can renew the domain name at any time before the expiration date or after the domain expires, as long as it falls within a grace period set by the registrar.

If a user acquires a block of IP addresses from an ISP, the ISP typically imposes a charge of IP space within the contract. If an IP block is acquired through ARIN, there is an annual fee associated with it. Such contractual information is not kept in the *whois* databases.

5-2 IP MANAGEMENT WITH DHCP

This section reviews the hierarchy of DHCP. You will understand how the BOOTP and DHCP processes work, as well as the steps for incorporating DHCP service into a campus network.

An IP address is one of the most basic pieces of information needed for a computer to communicate on a network. An IP address can be configured either manually or it can be assigned dynamically. In the manual process, a network administrator assigns an IP address to a user computer. Then, either the administrator or the user has to configure the computer's network settings with the assigned IP address along with other network parameters, such as the subnet mask, default gateway, domain name, and domain name servers. This can be a tedious process, especially when it involves multiple machines.

This process can be automated to some extent using a program called **BOOTP** for IP assignment. BOOTP stands for Bootstrap Protocol, and it enables computers to discover their own IP addresses. BOOTP uses UDP as its transport protocol. It utilizes UDP port 67 for BOOTP server (BOOTPS) and UDP port 68 for BOOTP client (BOOTPC). When a client requests an IP address, it is assigned to the Ethernet address (MAC address) based on the BOOTP record. In this case, the IP and MAC addresses have a one-to-one relationship. This means that a MAC address must be preregistered before its static or fixed IP address can be assigned. This makes BOOTP less robust for mobility and less efficient in IP space management.

Dynamic Host Configuration Protocol (**DHCP**) enhances BOOTP and simplifies the steps for IP assignment even further. DHCP is a superset of BOOTP and runs on the same UDP port numbers. Therefore, it interoperates with BOOTP clients. DHCP's function is to assign a pool of IP addresses to requesting clients. In this process, DHCP requests an IP address from the DHCP server. The DHCP server retrieves an available IP address from a pool of addresses dedicated to the subnet of the requesting client. The IP address is passed to the client, and the server specifies a length of time that the client can hold the address. This is called the **lease time**. This feature keeps an unused computer from unnecessarily tying up an IP address. DHCP servers will typically grant IP addresses for a limited time and the DHCP clients are responsible for renewing the address before it expires. Along with an IP address, other network settings such as gateway, subnet mask, DNS, Time Server, LDAP server, boot server, and boot filename can be included.

DHCP was originally designed to work on a local physical subnet, where both the DHCP server and the DHCP clients reside in the same LAN. When a computer is configured to obtain an IP address automatically or to use the DHCP option, the typical process of requesting an IP address with DHCP is as follows:

1. The client boots up and broadcasts a **DHCP Discover** message on its local network. This is a broadcast, meaning that the message is sent to all computers in the LAN.

BOOTP
Bootstrap Protocol. A network protocol used by a network client to obtain an IP address from a configuration server.

DHCP
Dynamic Host Configuration Protocol. The protocol used to assign a pool of IP addresses to requesting clients.

Lease Time
The amount of time that a client can hold an IP address.

DHCP Discover
This is a broadcast, meaning that the message is sent to all computers in the LAN.

DHCP Offer
The DHCP server listening on the LAN will take the packet, retrieve an available IP address from the address pool, and send the address to the client.

DHCP Request
Formally request and confirm the offered IP with the server.

DHCP ACK
A unicast packet sent back to the client with the same IP information.

2. A DHCP server listening on the LAN will take the packet, retrieve an available IP address from the address pool, and send the address to the client in a form of **DHCP Offer** message. The server sends the IP address and the server will send the lease time and other necessary network parameters, such as subnet mask, default gateway, domain name server, and so on. The DHCP Offer is a unicast message from the server to the client.
3. The client receives the OFFER from the server and agrees to use the lease. It replies back to the server in a form of **DHCP Request** to formally request and confirm the offered IP with the server. The DHCP Request is a broadcast message.
4. The server receives the REQUEST and sends back a **DHCP ACK**, which is a unicast packet, back to the client with the same IP information. The client applies the IP address and its network settings to the computer; then, it is ready to make network connections. An example of this process is provided in Figure 5-1.

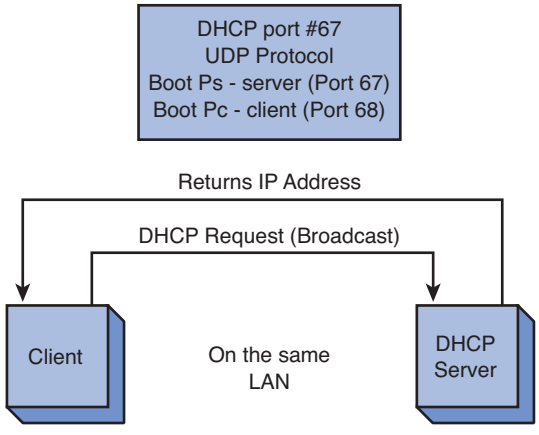

FIGURE 5-1 An example of a DHCP server and client in the same LAN

ipconfig/release
Command used to release the current IP address.

ipconfig/renew
Command used to initiate the DHCP process.

Automatic Private IP Addressing (APIPA)
A self-assigned IP address in the range of 169.254.1.0–169.254.254.255.

When a computer boots up, its network software will automatically engage in a DHCP process. However, this process can be invoked by using the command line. In Windows, the command **ipconfig /release** can be used to release the current IP address, then the command **ipconfig /renew** can be used to initiate the DHCP process. A frequently asked question is, "What happens if the DHCP server is not available?"

This scenario happens more often than one would think. A DHCP server could be offline or a broken network path to the server can cause the server to not be available. When this happens, a DHCP client will use a self-assigned IP address known as **Automatic Private IP addressing (APIPA)**. APIPA uses a reserved IP range of 169.254.1.0–169.254.254.255. The client will select a random IP address in that range with a netmask of 255.255.0.0. The client will then send a gratuitous ARP packet asking for that IP address to see if any other machine is using it. If there is an ARP reply by another machine, the client will generate another random IP address and try again. This could be deceiving to a typical user, because it looks as if the machine is getting an IP address; however, the machine will not be able to access the Internet.

DHCP was originally designed to only work in a local network. What if a DHCP server is on the other side of the router (for example, not in the same LAN)? The DHCP Discover and DHCP Request are broadcast messages, and they cannot propagate beyond their broadcast domains. Remember, routers don't pass broadcast addresses, so the DHCP broadcast is not forwarded. This situation requires that a DHCP relay be used, as shown in Figure 5-2. The DHCP relay sits on the same LAN as the client. It listens for DHCP requests and then takes the broadcast packet and issues a **unicast** packet to the network DHCP server. Unicast means that the packet is issued a fixed destination and therefore is no longer a broadcast packet. The DHCP relay puts its LAN address in the DHCP field so the DHCP server knows the subnet the request is coming from and can properly assign an IP address. The DHCP server retrieves an available IP address for the subnet and sends the address to the DHCP relay, which forwards it to the client.

Unicast
The packet has a fixed destination.

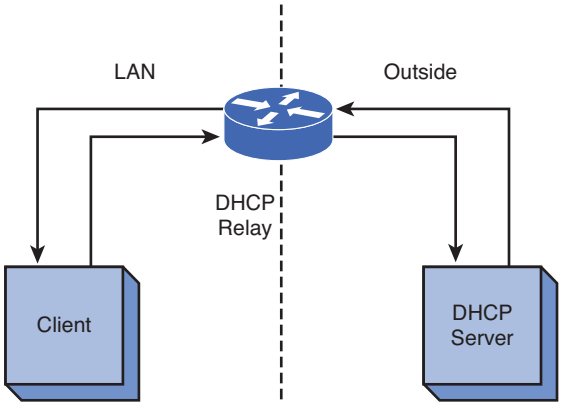

FIGURE 5-2 An example requiring the use of a DHCP relay

Cisco routers have a DHCP relay built in to their operating systems. The router command to enable the DHCP relay is Router(config-if)# **ip helper** [*ip address of the DHCP server*]. Note that this command is issued from the interface that connects to the LAN. In fact, the IP address for the interface is typically the gateway address for the LAN.

ip helper [*ip address of the DHCP server*]
The router command used to enable the router's DHCP relay function.

It was mentioned earlier that DHCP is a UDP protocol and uses port number 68 for the BOOTP-client and port 67 for the BOOTP-server. (BOOTP and DHCP use the same port numbers.) The BOOTP-client is the user requesting the DHCP service. The BOOTP-server is the DHCP server. The following discussion describes how these services are used in a DHCP request. The DHCP proxy on the router listens for the packets that are going to DHCP or BOOTP port numbers.

DHCP Data Packets

The following is a discussion on the packets transferred during a DHCP request. The network setup is the same as shown in Figure 5-2. The data traffic shown in this example will contain only the data packets seen by the client computer. The Wireshark protocol analyzer was used to capture the data packets. Figure 5-3 provides a portion of the captured data packets. Packet 31 is a DHCP broadcast

MT Discover
Message type discover, a DHCP Discover packet.

with a message type discover (**MT Discover**). In Wireshark, this is called the DHCP Discover packet. The destination for the packet is a broadcast, which is identified by the 255.255.255.255 destination IP address. The message source has a MAC address of Apple_f8:e2:dd, and the IP address is 0.0.0.0. The 0.0.0.0 indicates that an IP address has not been assigned to the computer. The source and destination ports are shown in the User Datagram Protocol (UDP) section in Figure 5-3. The source port is 68, which is for the Bootstrap Protocol Client (the computer requesting the IP address). The destination port is 67, the Bootstrap Protocol Server (the DHCP server).

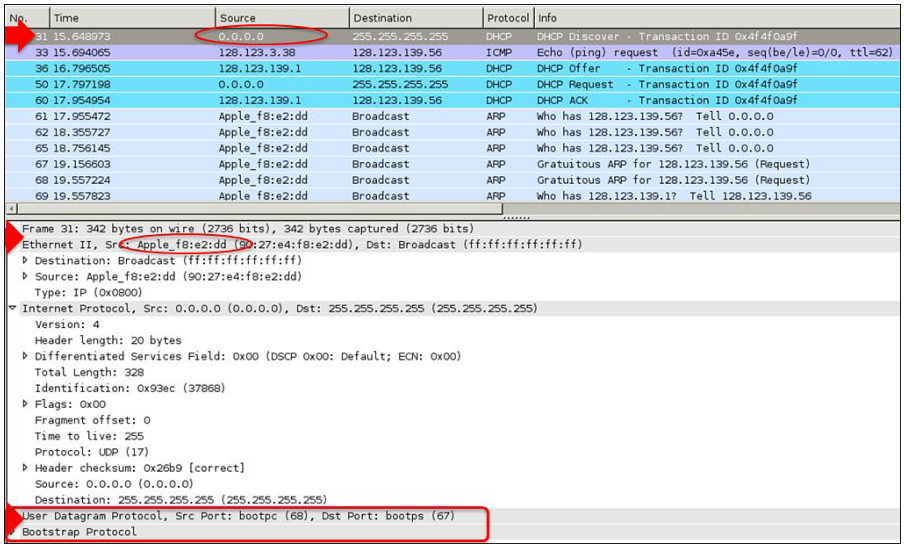

FIGURE 5-3 The captured DHCP packets

MT Offer
Message type offer, a DHCP offer packet.

Figure 5-4 shows a reply from the router acting as a relay agent on behalf of the DHCP server. Packet 36 is an offer of the IP address, 128.123.139.56 to the client. This is called the DHCP Offer packet (**MT Offer**). This packet contains the following information, as well as other network information that the client may need to connect to the network:

Domain name: (nmsu.edu)
DHCP server: (128.123.3.5)
Default gateway: (128.123.139.1)
Leased time: (12 hours)

MT Request
Message type request, a DHCP request packet.

Packet 50, shown in Figure 5-5, has a message type of **MT Request**, also called the DHCP Request packet. This packet is sent from the client back to the server that has been selected to provide the DHCP service (Note: It is possible for a campus LAN to have more than one DHCP server answering the DHCP request). Even though this packet is destined for the DHCP server as specified by the DHCP Server Identifier (128.123.3.5), the DHCP Request is still a broadcast message to its local network. The packet is picked up by the DHCP relay agent (router) and sent to the DHCP server. This means that the client is accepting the IP address offer.

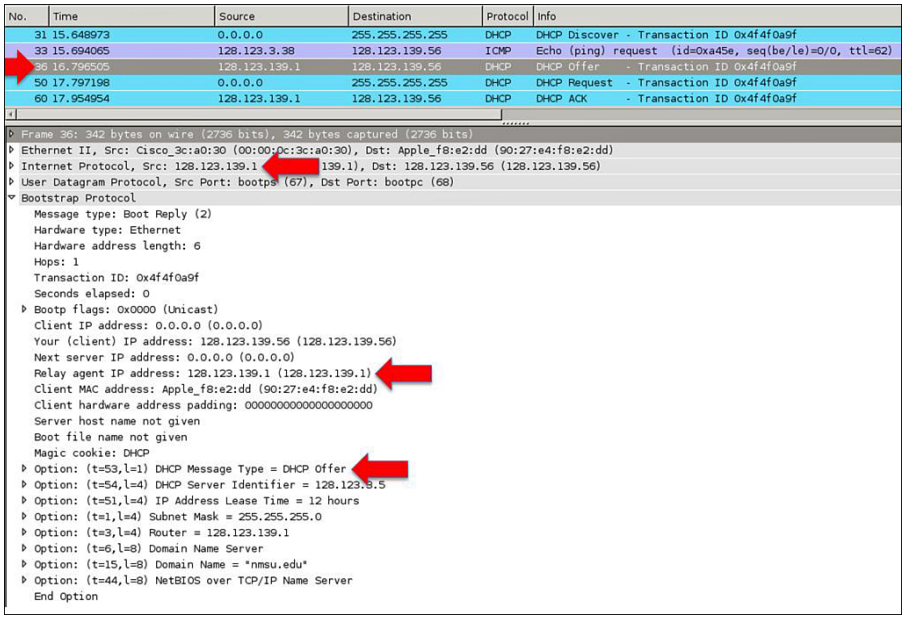

FIGURE 5-4 The DHCP offer packet

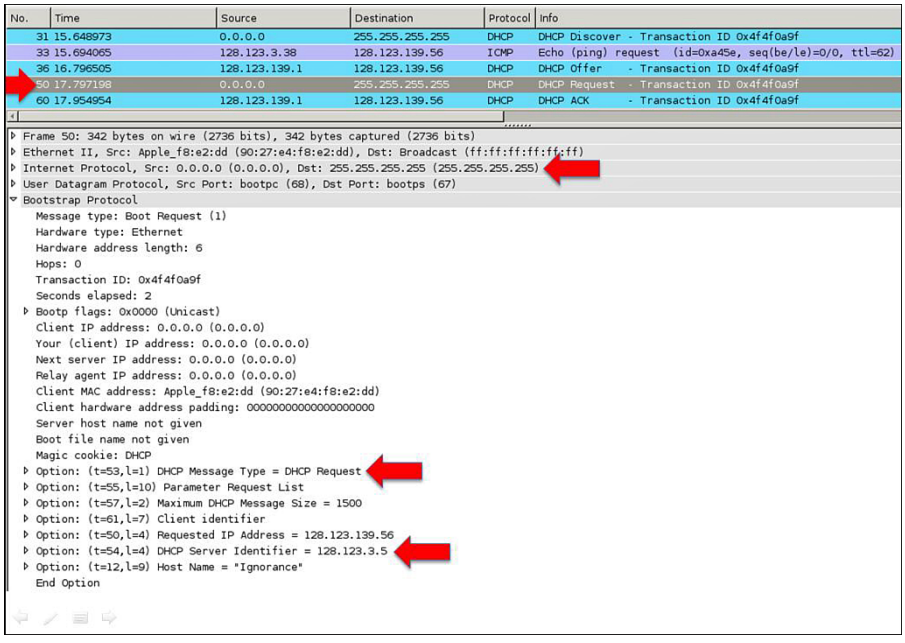

FIGURE 5-5 The DHCP Request packet

Packet 60, shown in Figure 5-6, is a message type of **MT ACK** or DHCP ACK. The DHCP server is acknowledging the client's acceptance of the IP address from the DHCP server. The packet is being relayed to the client via the router.

MT ACK
Message type acknowledgment, a DHCP ACK packet.

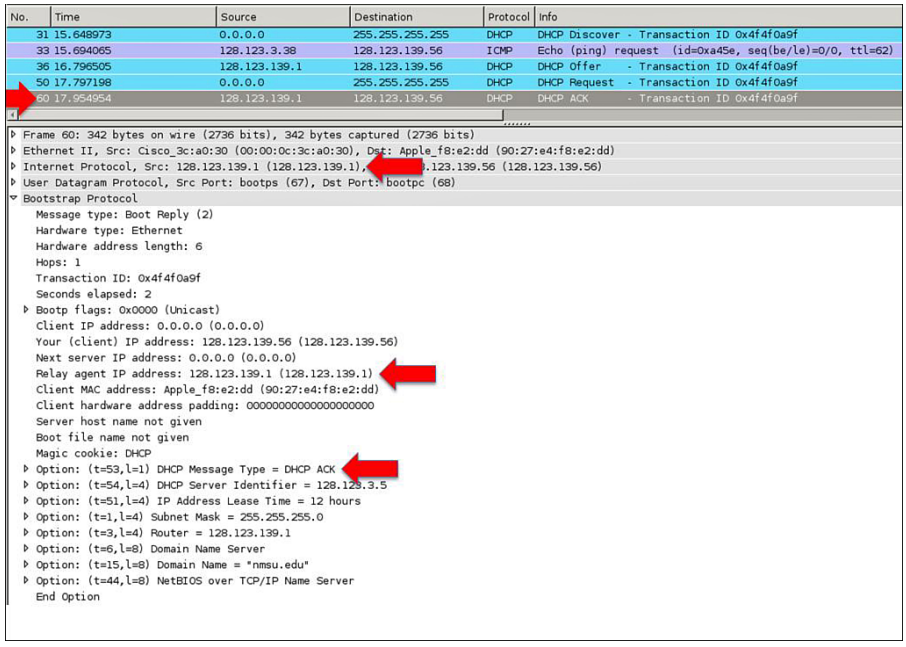

FIGURE 5-6 The DHCP ACK packet

ARP Broadcast
Used to inform everyone on the network that it now is the owner of the IP address.

The client computer now has an IP address assigned to it, but before the client can use it, the client will perform an ARP broadcast for the assigned IP address to verify no one else has been assigned this address. This step is shown in Figure 5-7. In this case, the ARP query is asking who has the IP address 128.123.139.56 and reply to 0.0.0.0. If there is no reply, the machine safely assumes that the IP address is good. It then performs a gratuitous **ARP broadcast**, informing everyone on the network that it now is the owner of the IP address. This is shown in Figure 5-8.

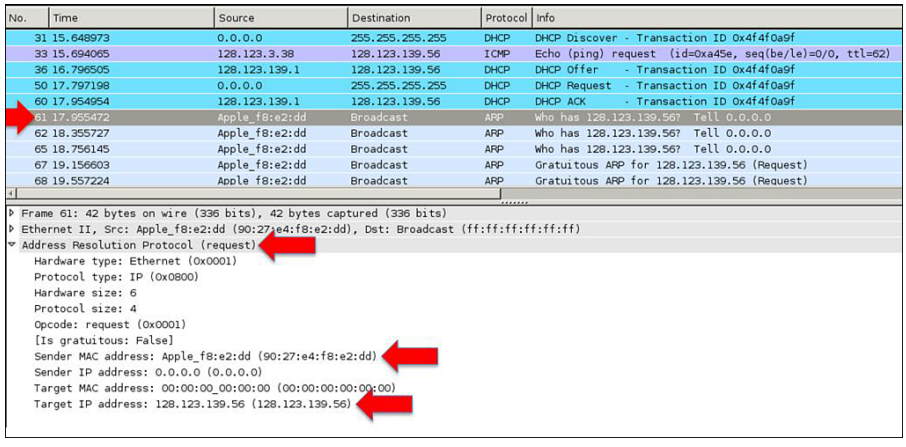

FIGURE 5-7 The ARP query asking who has the IP address 128.123.139.56

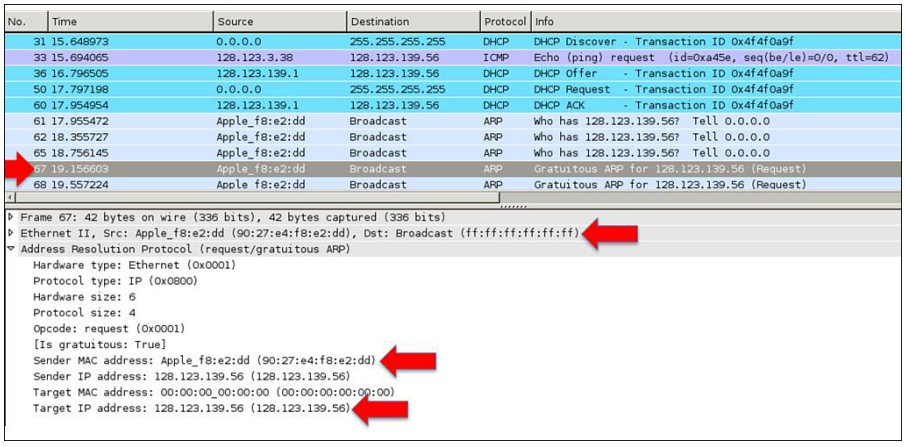

FIGURE 5-8 The gratuitous ARP broadcast for 128.123.139.56

DHCP Deployment

In a small office/home office (**SOHO**) environment, the network is typically small and only one router is needed. In this kind of network, a router performs simple routing functions, acts as a gateway to the outside world, and manages IP assignment via DHCP. Most network routers are capable of running the DHCP service, so it makes sense and is more cost-effective to deploy DHCP service at the router.

SOHO
Small office or home office network.

A DHCP configuration on any network device typically will start with defining a scope, a pool, or a range of the IP addresses that will be made available for allocation. Then, the network settings, such as gateway, subnet mask, DNS servers, domain name, leased time, and other information, will need to be assigned to the scope. The following is the example of how to configure a Cisco router to provide DHCP service.

When running DHCP service on a local router, the *ip helper-address* is not needed anymore to relay the DHCP information. This is because the DHCP service is in the same LAN as the client computer. The router must be placed in the router's configuration mode in order to configure the DHCP service. The DHCP pool must be configured with the command **ip dhcp** [*pool_name*]. Then, an IP network must be defined as the IP allocation pool. In this case, the IP addresses from 172.20.224.0–172.20.224.255 have been set aside for the address pool.

```
RouterA#conf t
Enter configuration commands, one per line. End with CNTL/Z.
RouterA(config)#
RouterA(config)# ip dhcp pool dhcp1pool
RouterA(dhcp-config)# network 172.20.224.0 255.255.255.0
```

Now that the basic DHCP pool is set up, it is time to associate other network settings to it. The following steps show how to define the DNS server, domain name,

and gateway to the DHCP pool. The subnet mask is already defined as part of the network.

```
RouterA(dhcp-config)# dns-server 172.20.224.8
RouterA(dhcp-config)# domain-name et477.com
RouterA(dhcp-config)# default-router 172.20.224.1
```

Even though an entire Class C network is being configured for the DHCP network, a portion of it might be reserved for something else, like static IP machines and servers. The command **ip dhcp** *excluded-address* is used to exclude a portion of IP addresses from being allocated to the DHCP devices. In this example, the IP addresses from 172.20.224.0 to 172.20.224.20 are excluded leaving the rest of network addresses available for DHCP allocation use:

```
RouterA(config)# ip dhcp excluded-address 172.20.224.0 172.20.224.20
```

To verify the DHCP pool status and information, the command **show ip dhcp pool** can be used:

```
RouterA#show ip dhcp pool

Pool dhcp1pool :
 Utilization mark (high/low)     : 100 / 0
 Subnet size (first/next)        : 0 / 0
 Total addresses                 : 254
 Leased addresses                : 5
 Pending event                   : none
 1 subnet is currently in the pool :
 Current index   IP address range                  Leased addresses
 172.20.224.132  172.20.224.1 - 172.20.224.254    5
```

Binding
An association of the IP address to the DHCP server.

Also, to see what IP addresses have been allocated by the DHCP server, the command **show ip dhcp binding** can be used. The **binding** is an association of the IP address to the DHCP server:

```
RouterA#show ip dhcp binding
Bindings from all pools not associated with VRF:
IP address      Client-ID/          Lease expiration        Type
172.20.224.12   0013.2126.9f2d      Feb 12 2012 02:17 PM    Automatic
172.20.224.47   0100.1143.bdda.03   Feb 12 2012 11:17 AM    Automatic
172.20.224.53   0100.0f1f.e82f.45   Feb 12 2012 04:07 AM    Automatic
172.20.224.62   0100.1143.11bb.ed   Feb 12 2012 12:04 PM    Automatic
172.20.224.114  0100.1143.11bc.5d   Feb 12 2012 11:33 AM    Automatic
```

In a larger and more complex environment where there are multiple networks and multiple routers, deploying DHCP service at the routers is not as simple. Having to manage a different DHCP service for each network on multiple routers can be tedious, time-consuming, and inefficient. This is where centralized DHCP service fares better.

This setup offers a centralized management, which scales better and is easier to support. A typical setup is to run a DHCP service program on a centralized server. With centralized DHCP service, the IP address assignment is typically tracked by

the network administrator or the network operations center (NOC). The tracking information can include more than the IP and MAC addresses, and the user information can also be included. This information can be kept in a central log file or in the database so that the administrator can troubleshoot network problems. For example, a machine could be causing network problems possibly due to hacked or corrupted software. The NOC needs to be able to track down the network problem(s). The NOC database will have the MAC address, the IP address, and the name of the person who uses the computer.

Also, in a large environment, DCHP pools are usually planned and pre-allocated. IP addresses are assigned by NOC based on where the subnet for the computer is located. The subnet could be in a building, a floor of the building, a department, and so on. The subnets are created by the network administrators based on the expected number of users (hosts) in a subnet. For example, the 192.168.12.0 network, shown in Figure 5-9, has been partitioned into four subnets. The network addresses for each of the subnets are provided in Table 5-1. Any computer in subnet B is assigned one of the 62 IP addresses from the range 192.168.12.65 to 192.168.12.126. Remember that the first IP address in the subnet is reserved for the network address and the last is reserved for the broadcast address.

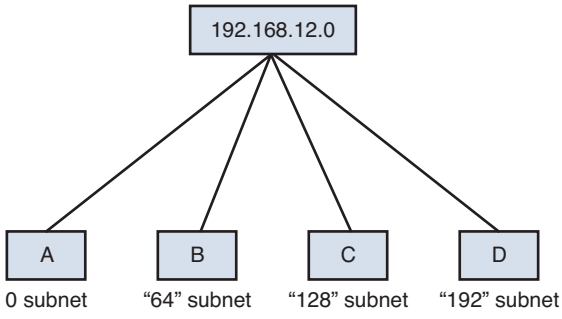

FIGURE 5-9 IP assignment of computers in a network’s subnet

TABLE 5-1 **Subnet Addresses for the Subnets Shown in Figure 5-9**

Subnet	Network Address	Broadcast Address	Subnet Mask
A	192.168.12.0	192.168.12.63	255.255.255.192
B	192.168.12.64	192.168.12.127	255.255.255.192
C	192.168.12.128	192.168.12.191	255.255.252.192
D	192.168.12.192	192.168.12.255	255.255.255.192

5-3 SCALING THE NETWORK WITH NAT AND PAT

It was mentioned in Chapter 1, "Network Infrastructure Design," that public IP addresses are a commodity, and not many institutions have a luxury of using public IP addresses in their network. In most cases, the demand of IP network devices exceeds the number of public IP addresses assigned to them. Most institutions have to use private IP addresses in their network. These private IP addresses must be able to communicate with outside or Internet hosts. This cannot be done because the private IP addresses are not routable on the Internet. These private IP addresses must be translated to public IP addresses using techniques like Network Address Translation (NAT) or Port Address Translation (PAT) for use on the Internet.

Network Address Translation (NAT)
A technique used to translate an internal private IP address to a public IP address.

Network Address Translation is a technique used to translate an internal private IP address to a public IP address before the packets leave the local network to the public network. NAT is typically implemented and deployed at the router facing the outside network. NAT is a one-to-one translation of a private IP address to a public IP address. This means that, for every connection made to the outside world, there must be a public IP address available. The public IP address is relinquished when it is no longer used or when the NAT timeout occurs. NAT is used not only for a way to communicate to the outside world; it can be used to hide the internal IP infrastructure of the network.

Port Address Translation (PAT)
A technique that uses the port number to identify the computer that established the Internet connection; also called many-to-one NAT and NAT overload.

To enhance NAT's limitation, **Port Address Translation (PAT)** was developed. PAT is sometimes referred to as many-to-one NAT and NAT overload, because of its capability to translate many IP addresses with a single public IP address or a handful of public IP addresses. PAT accomplishes this by using the TCP/UDP ports. The PAT process tracks a port number for the connection. The router stores the IP address and port number in a NAT lookup table. The port number differentiates the computer that is establishing a connection to the Internet because the router uses the same public IP address for all computers. This port number is used when a data packet is returned to the home network. The port number identifies the computer that established the Internet connection, and the router can deliver the data packet to the correct computer.

For example, if computer 1 establishes a connection to a website on the Internet, the data packets from the website are sent back to computer 1 using the network's routable public IP address. This first step enables the data packet to be routed back to the home network. Next, the router uses the NAT lookup table and port number to translate the destination for the data packet back to the computer 1 private IP address and original port number, which might be different. Table 5-2 demonstrates an example of a PAT table of a router. The router translates the private IP addresses to the public routable IP address assigned by the ISP. Additionally, the router tracks a port number with the public IP address to identify the computer. For example, the computer with the private IP address of 10.0.0.1 is assigned the public IP address 12.0.0.1:2000, where 2000 is the port number tracked by the router. The term NAT is used more generally than PAT and, most times, it covers PAT.

TABLE 5-2 **Example of a Router's PAT Table**

Inside IP	Inside Port	Outside IP	Outside Port
10.0.0.1	2000	12.0.0.1	2000
10.0.0.2	3000	12.0.0.1	3000
10.0.0.2	30001	12.0.0.1	4000
10.0.0.3	3000	12.0.0.1	5000
10.0.0.3	20010	12.0.0.1	6000

Configuring NAT

When dealing with NAT on Cisco routers, one must be familiar with the Cisco terminologies. Cisco uses the term ***local address*** to define any IP address that is on the inside of or internal to the network. The term ***global address*** is used to define any IP address that is on the outside of or external to the network. The first step of configuring NAT is to define the NAT points on a router's interfaces to designate the inside area and the outside area. Only when a packet passes through from the inside interface to the outside interface, a NAT will occur. The following example demonstrates how to define NAT points. The FastEthernet0/0 will be the inside interface and the FastEthernet0/1 will be the outside interface. This concept is graphically defined in Figure 5-10. To define the NAT area, the command **ip nat inside/outside** is used:

Local Address
Defines any IP address that is on the inside of or internal to the network.

Global Address
Defines any IP address that is on the outside of or external to the network.

```
RouterA#conf t
Enter configuration commands, one per line. End with CNTL/Z.
RouterA(config)#
RouterA(config)# interface FastEthernet0/0
RouterA(config-if)# ip nat inside
RouterA(config)# interface FastEthernet0/1
RouterA(config-if)# ip nat outside
```

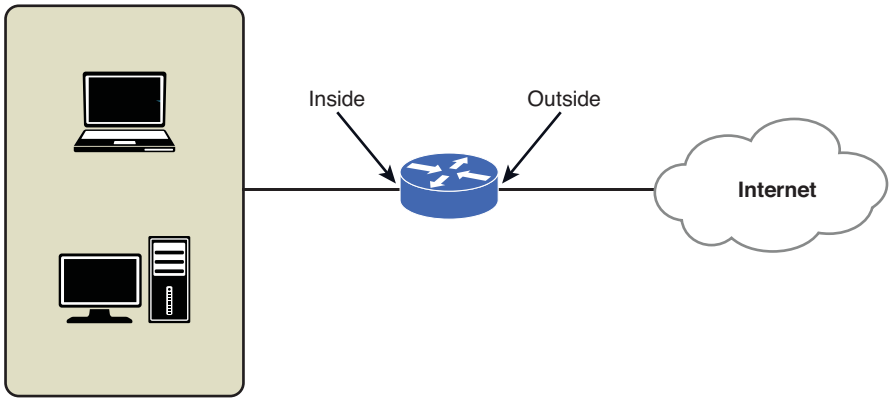

FIGURE 5-10 A graphical depiction of NAT inside and out

Static NAT
A fixed one-to-one mapping of an inside IP address to an outside IP address.

After the NAT interfaces are defined, a NAT statement can be configured. There are several types of NAT. One of them is static NAT. A **static NAT** is a fixed one-to-one mapping of an inside IP address to an outside IP address. The static NAT command is **ip nat inside source static** [*local_ip*] [*global_ip*]. The following example demonstrates how to configure a static NAT on a Cisco router:

```
RouterA(config)#
RouterA(config)# ip nat inside source static 10.0.0.5 12.0.0.5
```

The preceding command will entirely map the inside private IP address of 10.0.0.5 to the outside public IP address of 12.0.0.5. There is no port translation when the internal host 10.0.0.5 is making a connection outside. This host appears to the outside world as 12.0.0.5, and it is accessible from the outside via the very same public IP address. In the case where an internal server needs to be reached from the outside, then a static NAT can be used. However, this brings up a security concern of exposing an IP address entirely to the external network, so a static NAT is usually discouraged.

A better approach is to expose only network ports that need to be accessible by the external network. For example, if a web server is on the internal network and it must be made accessible to the external network, there is no need to map an entire IP address to it. Because a web server runs on specific TCP ports, like 80 for HTTP and 443 for HTTPS, these TCP ports can be made available via the static NAT instead. Sometimes, this technique is referred to as static PAT, port mapping, port forwarding, or port redirect. The following example demonstrates how to configure static NAT statements to map TCP port 80 of the inside host to TCP port 80 of the outside public IP address as well as TCP port 443 to cover all web traffic ports:

```
RouterA(config)#
RouterA(config)# ip nat inside source static tcp 10.0.0.5 80 12.0.0.5
80
RouterA(config)# ip nat inside source static tcp 10.0.0.5 443 12.0.0.5
443
```

Dynamic NAT
A one-to-one mapping from an available global pool.

Another type of NAT configuration is **dynamic NAT**. With dynamic NAT, it is not a fixed one-to-one IP mapping; it is a one-to-one mapping from an available global pool. A router is assigned a pool of IP addresses that contains global IP addresses, and every inside host that tries to access a public network will be given an IP from the global pool. In order to configure dynamic NAT on Cisco routers, a pool of global IP addresses and an access list containing allowable inside IP addresses must be defined. The command **ip nat pool** *[pool_name] [start_ip_address] [end_ip_address]* **netmask** *[subnet_mask]* is used to create a NAT pool. A standard or extended access list can be used to create a set of networks that is allowed to use the NAT pool. Then, a dynamic NAT statement can be configured to allow the specified IP addresses to be translated using the specified global IP addresses from the pool. The following example shows step-by-step configuration:

```
RouterA(config)#
RouterA(config)# ip nat pool global_ip 12.0.0.6 12.0.0.26 netmask
255.255.255.0
RouterA(config)# access-list 1 permit 10.0.0.0 0.0.0.255
RouterA(config)# ip nat inside source list 1 pool global_ip
```

With dynamic NAT, there is a limitation of the global pool being depleted, because of a limited number of global IP addresses available for the one-to-one IP mapping. Another type of NAT configuration is **NAT overload** or PAT and is supposed to solve that problem. Similar to dynamic NAT, a pool of global IP addresses and an access list containing allowable inside IP addresses must be defined. (Note: The use of access lists is discussed in Chapter 7, "Network Security"). Then, the PAT statement can be configured just by issuing the keyword **overload**. The following example shows step-by-step configuration. This example presents a slightly different way of doing the access list. Access-list 101 is an extended access list, but it still does the same as the access list 1 from the previous example, which is to allow only the network 10.0.0.0/8.

NAT Overload
Another name for PAT.

```
RouterA(config)#
RouterA(config)# ip nat pool global_ip 12.0.0.6 12.0.0.26 netmask
255.255.255.0
RouterA(config)# access-list 101 permit ip 10.0.0.0 0.0.0.255 any
RouterA(config)# ip nat inside source list 101 pool global_ip overload
```

If one cannot acquire a pool of global IP addresses from an ISP, then the NAT overload will have to be configured using the router interface facing the public network. This interface will have a public IP address assigned to it. For example, Router's FastEthernet 0/1 is the interface facing the outside public network. It has an IP address of 12.0.0.2. The maximum theoretical number of ports that a single IP can use is about 64,000. Even though the practical limit will not be as high due to a router's hardware limitation, a single public IP address will be more than enough to support a small to medium network's simultaneous connections to the external network. A NAT overload can then be configured to use the interface's IP address as the global address. In this case, the global IP pool is not needed, only an access list is required. The following example shows step-by-step configuration:

```
RouterA(config)#
RouterA(config)# access-list 101 permit ip 10.0.0.0 0.0.0.255 any
RouterA(config)# ip nat inside source list 101 interface
FastEthernet0/1 overload
```

To display active translations, the command **show ip nat translation** is used. The command will show the active NAT table in column format. The column "Pro" is the protocol (TCP, UDP, ICMP) being translated. The "Inside global" is the global IP address used by the inside IP address after the NAT process. The "Inside local" is the inside IP address. The "Outside local" corresponds to the destination IP address of the inside local before the NAT translation. The "Outside global" corresponds to the destination IP address of the inside global after the NAT translation. The following output shows the display of the active NAT translations. The example shows the global IP address of 12.0.0.2, which is an outside interface IP address of Router A:

```
RouterA# show ip nat translation
Pro Inside global         Inside local            Outside local
Outside global
tcp 12.0.0.2:57425    10.10.70.5:57425      74.125.227.20:80
74.125.227.20:80
```

```
tcp 12.0.0.2:57426    10.10.70.5:57426      74.125.227.17:80
74.125.227.17:80
tcp 12.0.0.2:57427    10.10.70.5:57427      173.194.69.94:80
173.194.69.94:80
tcp 12.0.0.2:53222    10.10.70.6:53222      192.65.78.152:80
192.65.78.152:80
tcp 12.0.0.2:53395    10.10.70.6:53395      192.65.78.150:80
192.65.78.150:80
tcp 12.0.0.2:53424    10.10.70.6:53424      216.38.172.205:1935
216.38.172.205:1935
tcp 12.0.0.2:54816    10.10.70.6:54816      74.125.224.175:80
74.125.224.175:80
tcp 12.0.0.2:55932    10.10.70.6:55932      192.65.78.151:80
192.65.78.151:80
tcp 12.0.0.2:57256    10.10.70.7:57256      74.125.224.173:80
74.125.224.173:80
tcp 12.0.0.2:58003    10.10.70.7:58003      66.220.145.45:80
66.220.145.45:80
tcp 12.0.0.2:59108    10.10.70.10:59108     192.65.78.152:80
192.65.78.152:80
tcp 12.0.0.2:59283    10.10.70.10:59283     96.7.191.139:80
96.7.191.139:80
tcp 12.0.0.2:59555    10.10.70.10:59555     96.7.191.139:443
96.7.191.139:443
tcp 12.0.0.2:1153     10.10.70.15:1153      65.55.223.34:40008
65.55.223.34:40008
tcp 12.0.0.2:1312     10.10.70.15:1312      192.65.78.152:80
192.65.78.152:80
tcp 12.0.0.2:1965     10.10.70.15:1965      213.146.189.204:12350
213.146.189.204:12350
tcp 12.0.0.2:1966     10.10.70.15:1966      213.146.189.205:12350
213.146.189.205:12350
udp 12.0.0.2:1882     10.10.70.16:1882      75.178.0.20:9967
75.178.0.20:9967
udp 12.0.0.2:1882     10.10.70.16:1882      85.27.9.65:15268
85.27.9.65:15268
udp 12.0.0.2:1882     10.10.70.16:1882      157.55.235.142:40019
157.55.235.142:40019
udp 12.0.0.2:1882     10.10.70.16:1882      157.56.52.14:40019
157.56.52.14:40019
udp 12.0.0.2:1882     10.10.70.16:1882      173.17.231.210:30850
173.17.231.210:30850
udp 12.0.0.2:1882     10.10.70.16:1882      178.207.64.233:7563
178.207.64.233:7563
udp 12.0.0.2:9001     10.10.70.16:9001      66.151.151.20:5062
66.151.151.20:5062
tcp 12.0.0.2:3593     10.10.70.16:49262     111.221.74.38:40008
111.221.74.38:40008
tcp 12.0.0.2:49307    10.10.70.16:49307     193.120.199.12:12350
193.120.199.12:12350
```

5-4 DOMAIN NAME SERVICE (DNS)

This section examines the Domain Name Service (DNS) services typically available in a campus network. Domain Name Service (**DNS**) translates a human-readable name to an IP address or an IP address to a domain name. The translation of a name to an IP address is called **forward DNS** lookup or forward DNS resolution, and translation of an IP address to a domain name is called **reverse DNS** lookup or reverse DNS resolution. DNS runs on UDP protocol port 53.

The DNS is a tree hierarchy. Everything in DNS starts at the "." servers, or generally called root servers, which are at the top of the hierarchy, as illustrated in Figure 5-10. The root servers are well-known IP addresses that have been programmed into DNS servers. When the DNS is installed on a server, a list of the root server's IP addresses is automatically configured in the DNS. A file containing the list of the most up-to-date root servers is available for the public, and it can be downloaded at the IANA's website. The file is known as the **Root Hints file (root.hints)**. According to IANA, there are currently 13 root servers distributed around the world operated by different independent entities. Each server is typically a cluster of servers spreading throughout different regions or countries. Table 5-3 shows the current list.

DNS
Domain Name Service.

Forward DNS
Translation of a name to an IP address.

Reverse DNS
Translation of an IP address to a name.

Root Hints File (root.hints)
A file containing the list of the most up-to-date root servers.

TABLE 5-3 **International Root Servers**

Hostname	IP Address	Manager
a.root-servers.net	198.41.0.4	VeriSign, Inc.
b.root-servers.net	192.228.79.201	University of Southern California (ISI)
c.root-servers.net	192.33.4.12	Cogent Communications
d.root-servers.net	128.8.10.90	University of Maryland
e.root-servers.net	192.203.230.10	NASA (Ames Research Center)
f.root-servers.net	192.5.5.241	Internet Systems Consortium, Inc.
g.root-servers.net	192.112.36.4	US Department of Defense (NIC)
h.root-servers.net	128.63.2.53	US Army (Research Lab)
i.root-servers.net	192.36.148.17	Netnod
j.root-servers.net	192.58.128.30	VeriSign, Inc.
k.root-servers.net	193.0.14.129	RIPE NCC
l.root-servers.net	199.7.83.42	ICANN
m.root-servers.net	202.12.27.33	WIDE Project

TLD
Top-level domain.

Country Domain
Usually, two letters, such as .us (United States) or .ca (Canada), that define the location of the domain server for that country. The next level below the TLD is the domain level and then the subdomain level. A domain is an Internet domain name that represents the organization or the entity. For example, et477.net and et477.com are Internet domains. A domain can be divided further into subdomains, like group1.et477.com and group2.et477.com.

The next hierarchical level from the root is the top-level domain. As discussed in Section 5-1, the top-level domains (TLD) registries are managed by IANA and ICANN. Examples of generic top-level domains (**TLD**) are as follows: .com, .net, .org, .edu, .mil, .gov, .us, .ca, .info, .biz, and .tv. **Country domain**s are usually defined by two letters, such as .us (United States) and .ca (Canada). The primary domain server for that domain has to exist in the same country; for example, the .us primary domain server is located in the United States. Figure 5-11 shows the top-level domains and their relationship to the subdomains and root servers. Examples of country-coded TLDs are as follows: .us, .uk, .au, .ru, .cn, .jp, .de, .ca, and so on.

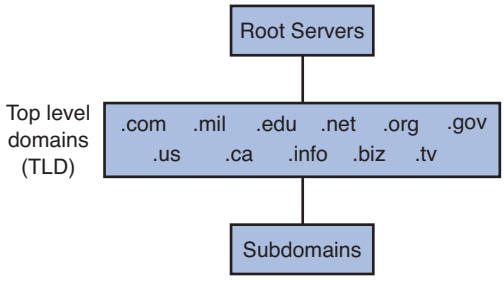

FIGURE 5-11 The domain name service tree hierarchy

DNS Tree Hierarchy

By having a tree hierarchy structure, DNS operates in a delegation mode starting from the root to subtree levels. This way, each level only has to maintain the information of the next level. With this structure, the root servers only know of the TLD servers, to which they can delegate the top level domain queries. The root servers will not know of **www.example.com**. They will delegate the query to a next level authoritative server, which repeats the delegation to the next level until it reaches the final destination that can authoritatively give the answer of **www.example.com**.

The following steps outline the typical name look-up process by a machine on the network (for example, if a machine called C1 wants to know the IP address for the www server at example.com):

1. C1 queries its network or campus domain name server as defined in its network settings for an IP address of the hostname www.example.com.
2. If the DNS server has the IP address of **www.example.com**, then it answers the name lookup query and the process is done. The DNS server could possess the IP address of **www.example.com**, because it is either an authoritative DNS server for the domain example.com or it has the IP information in its DNS cache.
3. If the campus or network DNS server does not have the IP address of **www.example.com**, it will start by querying one of the root servers to find the IP address of **www.example.com**.
4. The root server will return a list of delegated name servers for top level domain ".com" to the campus name server.

5. The campus DNS server will then query one of the .com TLD name servers on the list for the IP address of **www.example.com**.
6. A .com server will return a list of delegated servers for the domain "example.com" to the campus name server.
7. The campus DNS will then query one of the example.com domain name servers for the IP address of **www.example.com**.
8. A domain name server for example.com will return an IP address of **www.example.com** to the campus DNS server. When a domain name server can answer a query for that domain, it is said to be an authoritative DNS server of the domain.
9. The campus DNS server will update its DNS cache, so this multiple-step process of obtaining an IP address does not have to be repeated unnecessarily. Then, it will pass the IP information to the machine C1.

Dig (Domain Information Groper)
A DNS lookup utility.

nslookup
A DNS lookup utility.

The UNIX program called "dig" will next be used to illustrate the delegation process during the name lookup query. **Dig (Domain Information Groper)** is a DNS lookup utility that is readily available on most UNIX/Linux operating systems. Another widely used DNS lookup utility is **nslookup**, which comes with Windows, Mac, and UNIX/Linux. The command **nslookup** is a very simple command-line program; however, its output is not as detailed as that provided with dig. Throughout this section, we explore how to use both of the DNS utility programs: **dig** and **nslookup**. The following demonstrates the name lookup steps mentioned previously. The command **dig +trace www.example.com** is issued at the UNIX prompt. This command traces every step of the name lookup process.

```
[admin@noc ~]$  dig +trace www.example.com

; <<>> DiG 9.6-ESV-R4-P3 <<>> +trace  www.example.com
;; global options: +cmd
.                   492498      IN    NS    a.root-servers.net.
.                   492498      IN    NS    b.root-servers.net.
.                   492498      IN    NS    c.root-servers.net.
.                   492498      IN    NS    d.root-servers.net.
.                   492498      IN    NS    e.root-servers.net.
.                   492498      IN    NS    f.root-servers.net.
.                   492498      IN    NS    g.root-servers.net.
.                   492498      IN    NS    h.root-servers.net.
.                   492498      IN    NS    i.root-servers.net.
.                   492498      IN    NS    j.root-servers.net.
.                   492498      IN    NS    k.root-servers.net.
.                   492498      IN    NS    l.root-servers.net.
.                   492498      IN    NS    m.root-servers.net.
;; Received 512 bytes from 128.123.3.6#53(128.123.3.6) in 8 ms

com.                172800      IN    NS    d.gtld-servers.net.
com.                172800      IN    NS    c.gtld-servers.net.
com.                172800      IN    NS    l.gtld-servers.net.
com.                172800      IN    NS    b.gtld-servers.net.
```

```
com.                172800      IN    NS    e.gtld-servers.net.
com.                172800      IN    NS    i.gtld-servers.net.
com.                172800      IN    NS    k.gtld-servers.net.
com.                172800      IN    NS    h.gtld-servers.net.
com.                172800      IN    NS    m.gtld-servers.net.
com.                172800      IN    NS    a.gtld-servers.net.
com.                172800      IN    NS    j.gtld-servers.net.
com.                172800      IN    NS    g.gtld-servers.net.
com.                172800      IN    NS    f.gtld-servers.net.
;; Received 493 bytes from 192.228.79.201#53(b.root-servers.net) in 28
ms

example.com.            172800      IN    NS    a.iana-servers.net.
example.com.            172800      IN    NS    b.iana-servers.net.
;; Received 169 bytes from 192.54.112.30#53(h.gtld-servers.net) in 185
ms

www.example.com.        172800      IN    A     192.0.43.10
example.com.            172800      IN    NS    b.iana-servers.net.
example.com.            172800      IN    NS    a.iana-servers.net.
;; Received 97 bytes from 199.43.132.53#53(a.iana-servers.net) in 26
ms
```

The **dig** result shows that the first step for looking up the name **www.example.com**, the campus DNS server, 128.123.3.6, is queried on port 53, even though it is not specified as UDP port 53. The campus DNS server has a list of 13 root servers, as shown as **a – m.root-servers.net**. It queries one of the root servers, which happens to be **b.root-servers.net**. The b.root-servers.net returns a list of the .com TLD name servers, **a – m.gtld-servers.net**.

The **h.gtld-servers.net** is then queried by the campus DNS server for **www.example.com**. It replies back with a list of two name servers for the domain example.com, which are a.iana-servers.net and b.iana-servers.net. The campus DNS server then queries the name **a.iana-servers.net** for **www.example.com**, which it replies back with the IP address of 192.0.43.10.

Authoritative Name Server

A name server that is authorized and configured to answer DNS queries for a particular domain or zone.

The server **a.iana-servers.net** is an authoritative name server for the domain example.com. An **authoritative name server** is a name server that is authorized and configured to answer DNS queries for a particular domain or zone. So, how does one become an authoritative name server of a domain? The answer is that this is done when a domain is registered. To successfully register an Internet domain name, the authoritative name servers must be specified. So, the authoritative name servers of a domain can be found by using the **whois** command. In Section 5-1, the command **whois example.com** was issued, and it showed the registration information of the domain example.com, which included its name servers: A.IANA-SERVERS.NET and B.IANA-SERVERS.NET. These name servers are the true authoritative domain name servers of the domain example.com.

In the example, the query was done against the campus DNS server, 128.123.3.6. This DNS server is said to be a non-authoritative name server, because it does not contain a copy of the domain example.com; therefore, it is not authorized to answer the query. A non-authoritative name server will always query the authoritative name servers of the domain for the answer. A name lookup answer received by a client via a non-authoritative server is called a **non-authoritative answer**.

Non-Authoritative Answer
A name lookup answer received by a client via a non-authoritative server.

The next examples demonstrate the difference when querying the authoritative server and the non-authoritative server. The program **nslookup** is used in this demonstration. During the first case, **nslookp** will query against an authoritative server, such as the example.com domain and b.iana-servers.net. The second time, it will query against a non-authoritative server, such as a campus name server. Both will yield the same IP address result, but the second attempt will show that the answer is a non-authoritative, because it is coming from the campus DNS server.

```
authoritative server
[admin@noc ~]$  nslookup
> server b.iana-servers.net
Default server: b.iana-servers.net
Address: 199.43.133.53#53
> www.example.com
Server:     b.iana-servers.net
Address:    199.43.133.53#53

Name:    www.example.com
Address: 192.0.43.10
>
> server 128.123.3.6
Default server: 128.123.3.6
Address: 128.123.3.6#53
> www.example.com
Server:     128.123.3.6
Address:    128.123.3.6#53
Non-authoritative server:
Name:    www.example.com
Address: 192.0.43.10
>
```

In the preceding examples, the hostname **www.example.com** is said to be a fully qualified domain name (**FQDN**) because it contains a full path of the domain name. Not all DNS queries are done with an FQDN, because of a domain name suffix that is configured as part of the network settings configuration. For example, a domain suffix example.com may be configured as a default domain or a search domain on a host. Therefore, instead of issuing an FQDN of **www.example.com** or web.example.com, partial qualified domain names (**PQDN**) of www or web can be used. PQDNs serve as shorthand, so that users don't have to provide the full name of the host. This is a general practice in a campus network.

FQDN
Fully qualified domain name.

PQDN
Partial qualified domain name.

DNS Resource Records

RR
Resource record.

As mentioned previously, an authoritative name server is a name server that is authorized and configured to answer DNS queries for a particular domain or zone. The authoritative name server is in charge of managing the information about that zone or domain. The information of the domain and its hosts and services are defined by resource records (**RR**) and are organized by zones. The terms *domain* and *zone* are often used interchangeably.

When a name server is hosting a single domain with no subdomains, the domain and the zone are the same. However, this creates multiple zones within a domain when there are subdomains involved. For example, the domain example.edu could have subdomains of engineering.example.edu and business.example.edu. This creates three different zones: one for example.edu, one for engineering.example.edu, and one for business.example.edu. Each zone contains resource records that define or describe a domain or subdomain. The following are the common resource records that would accompany a zone.

SOA
Start of Authority.

SOA Resource Records **SOA** or Start of Authority is a mandatory RR for each zone. It marks the start of the zone and provides the technical details of the zone, such as zone name, the primary authoritative name server, the email address of the domain administrator, serial number of the domain, TTL (Time to Live) of the domain, refresh, retry, and expire time for the slave name server. The SOA record of the domain can be found by using either the **dig** or **nslookup** command. The following example uses the **nslookup** command to find the SOA record of the domain example.com.

```
C:\nslookup -query=SOA example.com
 Server:  192.168.1.1
 Address:  192.168.1.1#53

Non-authoritative answer:
example.com
        primary name server = dns1.icann.org
        responsible mail addr = hostmaster.icann.org
        serial  = 2011063168
        refresh = 7200 (2 hours)
        retry   = 3600 (1 hour)
        expire  = 1209600 (14 days)
        default TTL = 3600 (1 hour)
```

This shows that the IP address is 192.168.1.1 using UDP port 53. Notice that this states that this is a non-authoritative answer. Also, the technical details of the zone are listed such as the primary name server and responsible mail address.

A Record (Address Record)
This maps a hostname to an IP address.

A Resource Records The **A record** or Address record is the most common record in DNS. It is a hostname mapping to an IP address. For example, the host1 entry in domain network-B.edu is an A record. The A record is used by a DNS server at the parent company for network-B to convert the name host1.network-B.edu to an IP address. By default, both **dig** and **nslookup** command will yield the IP address of the hostname. It is not necessary to specify the option for A record when

using dig or nslookup. The following example uses the **nslookp** command to find the A record of the host **www.example.com**.

```
C:\nslookup  www.example.com
 Server:  192.168.1.1
 Address:  192.168.1.1#53

Non-authoritative answer:
Name: www.example.com
Address: 192.0.43.10
```

This shows that the IP address for **www.example.com** is 192.0.43.10.

PTR Resource Records The **PTR record** or Pointer record is a reverse of an A record. It is a mapping of an IP address to a hostname. It is sometimes referred to as a reverse record. The following example uses the **nslookp** command to find the PTR record or the reverse record of the IP address given as the result of www.example.com:

PTR Record (Pointer Record)
The reverse of an A record.

```
C:\nslookup -query=PTR 192.0.43.10
 Server:  192.168.1.1
 Address:  192.168.1.1#53

Non-authoritative answer:
10.43.0.192.in-addr.arpa name = 43-10.any.icann.org.
```

Interestingly, the answer is not the exact reverse hostname that one would expect. One would think that name result would be reciprocal to the previous exercise, which is www.example.com. In this case, the name listed is 43-10.any.icann.org instead. This is common on the Internet. This has to do with how the Internet domain name is registered and how the IP address is acquired. As we have learned, the Internet domain name can come from different sources. An Internet domain name is generally purchased from a domain registrar, and the IP address has to be allocated from ARIN (in North America) or from an ISP. This results in one entity being in charge of the forward DNS zone and another entity being in charge of the reverse DNS zone. The information is not usually synchronized, hence the result above.

CNAME Resource Records **CNAME (Canonical name) record** is generally called an alias. It allows another name to be defined and points to the real name. CNAME record is mapped to an A record. Similar to an A record query, it is not necessary to specify the option for CNAME record when using dig or nslookup. Both commands yield a canonical name or an alias of a hostname, if it exists. The following example reveals that www.iana.org is actually a name of an A record of ianawww.vip.icann.org:

CNAME (Canonical Name) Record
Generally called an alias of a hostname.

```
C:\nslookup  www.iana.org
 Server:  192.168.1.1
 Address:  192.168.1.1#53

Non-authoritative answer:
www.iana.org     canonical name = ianawww.vip.icann.org.
```

```
Name:  ianawww.vip.icann.org
Address: 192.0.32.8
```

CNAME is useful in applications like virtual web hosting. Being able to create multiple names and mapping these names to one canonical name that in turn associates to one IP address, allows multiple websites to be served by one server. The effect is seamless to general users. As a matter of fact, this is how most of the virtual services and cloud services are able to provide their services. For example, what if the website www.example.com moved to a cloud service provider by creating a CNAME record and then mapped the website to a specific cloud service provider's server? To general users, the website is still www.example.com, but the destination where the service is hosted is now different.

NS Record
Specifies the name of the authoritative name server of the domain.

NS Resource Records **NS record** or Name Server record is another mandatory RR for a zone. This specifies the name of the authoritative name server of the domain. The record must map to a valid A record, not an IP address or a CNAME. The NS records are associated with the domain, not a particular host. Therefore, one will need to look up the name server information based on the domain. The following example demonstrates the use of the **nslookup** command to lookup the NS records of the domain example.com.

```
C:\nslookup  -query=NS example.com
 Server:  192.168.1.1
 Address:  192.168.1.1#53

Non-authoritative answer:
example.com nameserver = a.iana-servers.net.
example.com nameserver = b.iana-servers.net.
```

In this case, the authoritative name servers are a.iana-servers.net and b.iana-servers.net.

MX Record
Specifies the email handling server of the domain.

MX Resource Records The **MX record** or Mail Exchange record specifies the email handling server of the domain. This is the server where all the incoming emails to the domain will go to. The MX record must also map to a valid A record, not an IP address or a CNAME. The MX record is a crucial piece of information in today's Internet. Without correct MX records, emails to the domain will stop flowing. The following example demonstrates the use of the **nslookup** command to search the MX records information of the domain network-b.edu. The command is issued per entire domain, just like the NS information.

```
C:\nslookup  -query=MX network-b.edu
 Server:  192.168.1.1
 Address:  192.168.1.1#53

 Non-authoritative answer:
network-b.edu  mail exchanger = 20 mx02.cloud.example.com.
network-b.edu  mail exchanger = 30 mx03.cloud.example.com.
network-b.edu  mail exchanger = 10 mx01.cloud.example.com.
```

The MX records yields three email servers for the domain network-b.edu. Each server has a different preference number, where 10 is the preference number of mx01.cloud.example.com server, 20 is the preference number of mx02.cloud.example.com server, and 30 is the preference number of mx03.cloud.example.com server. The preference number is sometimes referred to as the distance or the priority. The lowest preference number signifies the most preferred server. Therefore, the server mx01.cloud.example.com, which has the preference number of 10, is the most preferred mail server of the three. Email service is another popular service offered by many cloud service providers. Many entities do not have resources to manage the amount of emails going to and coming from their domain. Similar to how CNAME records are used to map to the cloud server, the MX records can be used in similar fashion. To move the email service of the domain to the cloud service, the MX will need to be changed to the cloud service provider's email servers. As shown in the preceding example, the domain network-b.edu is using cloud.example.com as their email cloud service.

TXT Resource Records The **TXT record** or Text record is used to hold arbitrary text information of the domain. Besides storing arbitrary information or comments for the domain, this record is being used increasingly more to validate the authenticity of the domain. One of its popular applications is to authenticate the email sender domain. **SPF** or Sender Policy Framework can be entered into a TXT record. This piece of information can be used as a validation of the legitimate sources of email from a domain. Another application is for the cloud service providers to validate the authenticity of the domain ownership. Many cloud service providers will ask for the proof of the domain ownership by providing the domain owner with a token value that needs to be added to the TXT record. The following example shows the TXT record with the specific token value (t=) for a cloud service.

TXT Record
Used to hold arbitrary text information of the domain.

SPF
Sender Policy Framework.

```
      C:\nslookup  -query=TXT network-b.edu
       Server:  192.168.1.1
       Address:  192.168.1.1#53

      Non-authoritative answer:
alumni.nmsu.edu  text = "v=msv1 t=3b6735dd2923c44e99c313ac4adb65"
```

SRV Resource Records The **SRV record** or Service record is used to identify a host or hosts that offer a specific type of service. This is sometimes called a service location record. The uniqueness of this type of record is its syntax. The SRV record has a syntax of _service._protocol.name (for example, _ldap._tcp.network-b.edu or _http._tcp.example.com). Not only does the SRV record provide typical host information, it also provides the TCP or UDP port of the service. The SRV record is used all the time with Microsoft Windows, especially in the Active Directory (AD) environment. The following example shows the SRV record of _ldap._tcp.network-b.edu. This service record allows a client to locate a server that is running the LDAP service for the domain network-b.edu.

SRV Record
Used to identify a host or hosts that offer that specific type of service.

```
      C:\nslookup  -query=SRV _ldap._tcp.network-b.edu
       Server:  192.168.1.1
       Address:  192.168.1.1#53

Non-authoritative answer:
```

```
_ldap._tcp.network-b.edu  service = 0 100 389 dc2.network-b.edu.
_ldap._tcp.network-b.edu  service = 0 100 389 dc1.network-b.edu.
```

The output yields two servers that can provide the service. Both of them have priority of 0, which is the highest, and the weight value of 100. Both of them provide the LDAP service on TCP port 389.

Administering the Local DNS Server—A Campus Network Example The primary records are the A records of a campus network and, as previously mentioned, are the most common records in DNS. These records contain the hostname and IP addresses for the computers. For example, network-B.edu has an assigned IP address of 172.16.12.1:

1. When a host pings www.network-B.edu, the host computer first checks its DNS cache; assuming the DNS cache is empty, the host then sends a DNS request to the campus DNS server. Typically, the host will know the IP addresses of the primary and secondary DNS server through either static input or dynamic assignment.
2. The request is sent to the primary DNS server requesting the IP address for www.network-B.edu. The primary DNS server is the authority for network-B.edu and knows the IP address of the hosts in the network.
3. The primary DNS server returns the IP address of www.network-B.edu, and then the ICMP process associated with a ping is started.

You might ask, "How does a PC in the campus network become part of the campus domain?" Specifically, how is an A record entered into the campus domain? Recall that the A record provides a host to IP address translation. Adding the PC to the campus domain is done either manually or dynamically.

The Steps for Manually Adding a Client to the Campus Network The steps for manually updating the DNS A records are graphically shown in Figure 5-12 and are as follows:

1. A client PC updates the A record when an IP address is requested for a computer.
2. The user obtains the PC name and the PC's MAC address.
3. This information is sent to the network administrator or the NOC.
4. The NOC issues an IP address to the client, updates the NOC database of clients on the network, and enters a new A record into the primary DNS. The entry is made only on the primary DNS.
5. The entry will be later replicated on the secondary DNS.

FIGURE 5-12 Manually updating the A record

The Steps for Dynamically Adding a Client to the Campus Network A new A record can be entered dynamically when the client computer obtains an IP address through DHCP registration. This is graphically depicted in Figure 5-13. The DHCP server will issue an IP address to the client and at the same time send an updated A record to the network's primary DNS. This process is called dynamic DNS (DDNS) update. The client name and the IP and MAC addresses are stored in the DHCP database.

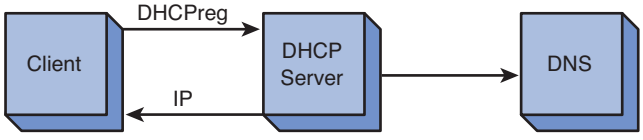

FIGURE 5-13 Dynamic updating of the A record using DHCP

Why obtain the MAC address when entering the information into DNS? This record is used to keep track of all the machines operating on the network. The MAC address is a unique identifier for each machine. The MAC address is also used by BOOTP, which is a predecessor to DHCP. This is where a MAC address is specifically assigned to one IP address in the network.

Reverse DNS returns a hostname for an IP address. This is used for security purposes to verify that your domain is allowed to connect to a service. For example, pc-salsa1-1 (10.10.20.1) connects to an FTP server that allows only machines in the salsa domain to make the connection. When the connection is made, the FTP server knows only the IP address of the machine making the connection (10.10.20.1). The server will use the IP address to request the name assigned to that IP. A connection is made to the salsa domain server, and the salsa DNS server returns pc-salsa1-1 as the machine assigned to 10.10.20.1. The FTP server recognizes this is a salsa domain machine and authorizes the connection.

SUMMARY

This chapter provided a detailed look at configuring and managing the network infrastructure. The concept of the domain name and IP assignment was first examined. The whois protocol was demonstrated and we showed how the databases can be queried to get information of an Internet domain name and an IP space. The issues of IP management with DHCP were also examined and how a name is translated to an IP address. Public IP addresses are a commodity and the steps for using NAT/PAT to translate a private IP address to a public IP address were demonstrated. The benefit is the network can use private IP address assignments inside and only require a limited number of shared public IP addresses for outside access. The DNS services typically available in a campus network have been presented. The DNS Tree Hierarchy was examined and examples of using the **dig** and **nslookup** commands were presented. The concept of an authoritative and non-authoritative name server was presented. A definition of the DNS source records was also presented.

QUESTIONS AND PROBLEMS

Section 5-1

1. What are the two key elements used by the general population when accessing websites on the Internet?
2. What is the purpose of IANA?
3. COM is an example of which of the following?
    a. DNS root zone for a generic (cc) top-level domain
    b. DNS root zone for a generic (d) top-level domain
    c. DNS root zone for a generic (g) top-level domain
    d. DNS root zone for a generic (int) top-level domain
4. gTLD stands for which of the following?
    a. Global top-level domain
    b. Generic top-level domain
    c. Gated top-level domain
    d. None of the above
5. An example of a ccTLD is which of the following? (Select two.)
    a. .net
    b. .uk
    c. .org
    d. .au

6. What is the .int domain registries?
7. What is the purpose of the IDN (Internationalized Domain Name) practices repository?
8. What are the three primary functions of IANA? (Select three.)
    a. Domain Name Management
    b. Portal Assignments
    c. Numbers Resource Management
    d. Protocol Assignments
    e. ASE Number Allocation
9. What organization is responsible for IP address assignment in North America?
10. ARIN's responsibility is to assign IP addresses to which of the following? (Select two.)
    a. Internet service providers
    b. Home networks
    c. Corporate networks
    d. Large end users
11. Which of the following are world Regional Internet Registries?
    a. AfriNIC
    b. AFrNIC
    c. LACNIC
    d. ARIN
    e. AIRN
12. What is the purpose of the in-addr.arpa domain?
13. Who handles the assignment of a domain name?
    a. ICANN
    b. Domain registrar
    c. Network administrator
    d. TLD
14. What protocol is used to query databases that store user registration information of an Internet domain name and IP space?
    a. whois protocol
    b. whereis protocol
    c. whatis protocol
    d. None of the above

Section 5-2

15. In regards to campus DHCP service, the IP address assignment is based on what?
16. How are BOOTP and DHCP related?
17. Define lease time.
18. What networking function is required if the DHCP server is not on the same LAN? Why is this networking function required?
19. What command enables a DHCP relay on a Cisco router?
20. What are the port numbers for the DHCP protocol?
21. What command is used to release a current IP address on a computer?
22. What command is used to initiate the DHCP process?
23. What happens if a DHCP server is not available?
    a. The client will issue a global broadcast to search for available DHCP server.
    b. The host computer will issue unicast packets to the 169.254.1.1 address and then obtain an IP address.
    c. A DHCP client will use a self-assigned IP address known as Automatic Private IP addressing (APIPA).
    d. The ipconfig/redo command is automatically issued to establish connectivity.
24. What is the command on the router to enable the DHCP relay function?
25. What information is contained in the MT offer (DHCP Offer) packet? (Select three.)
    a. Default gateway
    b. Leased time
    c. Internet address
    d. IP address of the Domain Name server
    e. Hostname
26. What is a gratuitous ARP broadcast?
27. What is the purpose of the following command?

```
ip dhcp pool address-pool
```

28. The following commands are entered into a router. Explain what this does.

```
RouterA(dhcp-config)# dns-server 192.168.10.52
RouterA(dhcp-config)# domain-name networks.com
RouterA(dhcp-config)# default-router 192.168.10.1
```

29. What information does NOC typically associate with an IP address?

Section 5-3

30. How is IP addressing typically handled in a home network?
31. What is Port Address Translation (PAT)?
32. A router on a home network is assigned an IP address of 128.123.45.67. A computer in the home network is assigned a private IP address of 192.168.10.62. This computer is assigned the public IP address 128.123.45.67:1922. Which IP address is used for routing data packets on the Internet? Is overloading being used?
33. For Cisco routers, what is a local address?
34. For Cisco routers, what is a *global address?*
35. If an interface Fastethernet0/0 is the interface for the internal private LAN and the interface Fastethernet0/1 is the interface facing the public Internet, provide the configuration for the appropriate NAT interfaces.
36. What does the following statement do?

```
RouterA(config)# ip nat inside source static 10.10.20.1
128.123.14.10
```

37. What does the following command do?

```
RouterA(config)# ip nat inside source static tcp 192.168.12.5 443
12.0.0.5 443
```

38. What does the following command do? Does this bring up any security concerns?

```
RouterA(config)# ip nat inside source static 10.10.20.1 15.1.1.2
```

39. The command **show ip nat translation** is entered on a router. The following information is displayed. What is this showing?

```
RouterA# show ip nat translation
Pro Inside global   Inside local      Outside local
Outside global
tcp 15.1.1.2:35425  10.10.20.1:35425  55.105.35.15:80
55.105.35.15:80
```

40. What is the maximum theoretical number of ports that a single IP can use?
41. What is dynamic NAT?

Section 5-4

42. List 11 top-level domains.
43. What is the purpose of a root server in DNS?
44. A new network wants to obtain a domain name. The first step is what?
45. The hostname and IP address for a computer is stored in what for a campus DNS service?
46. How is it possible for the command **ping www.networkB.edu** to find the destination without an IP address?
47. What is the purpose of reverse DNS? Where is it used?

48. What is the purpose of the reverse domain name service?
49. What is the purpose of the root hints file?
50. In regards to the Internet, what is a domain?
51. What is an authoritative name server?
52. Explain how a machine obtains the IP address of a website on the Internet.
53. The following entry is made on a UNIX server. Describe what this is doing.

```
[admin@noc ~]$  dig +trace www.example.com
```

54. What does it mean to be a non-authoritative name server?
55. What is a fully qualified domain name (FQDN)?
56. What is the difference in a domain and a zone?
57. What is the Start of Authority?
58. This record is a mapping of an IP address to a hostname. It is sometime referred to as a reverse record.
59. What is this information showing?

```
www.iana.orgcanonical name = ianawww.vip.icann.org.
```

60. Where is the authoritative name server of the domain listed?
61. What is the purpose of the TXT record?

Critical Thinking

62. Describe the typical process for requesting an IP address using DHCP.
63. How does NAT (Network Address Translation) help protect outsider access to computers in the home network?
64. Why would the pointer record not be the exact reverse of the A record?
65. The following query is submitted. What is this information showing? Which server is the preferred server? What is the cloud being used for?

```
C:\nslookup  -query=MX network-b.edu
    Server:  192.168.1.1
    Address:  192.168.1.1#53

   Non-authoritative answer:
et477.com  mail exchanger = 15 mx02.cloud.sample.com.
et477.com  mail exchanger = 25 mx03.cloud.sample.com.
et477.com  mail exchanger = 5 mx01.cloud.sample.com.
```

This page intentionally left blank