

**Chapter
2**

**TELECOMMUNICATIONS
AND WINDOWS NT**

*Get on the
Fast Track!*




TM

**SYS-ED/
COMPUTER
EDUCATION
TECHNIQUES, INC.**

Objectives

You will learn:

- C Subnets, subnet masks, gateways, and routers.
- C TCP/IP.
- C Routing protocols.
- C How to map from domain names to IP addresses.
- C The role and function of DHCP, WINS, and DNS.



1 TCP/IP Overview

TCP/IP stands for Transmission Control Protocol/Internet Protocol.

TCP/IP is a suite of protocols in which each protocol has its own specialized function. Virtually every type of computer and operating system has available and supports TCP/IP as a networking protocol.

IPX/SPX and NetBEUI are proprietary protocols developed and ultimately controlled by an individual vendor. Unlike IPX/SPX or NetBEUI, TCP/IP is controlled by no company. TCP/IP is an open standard controlled by the Internet Engineering Task Force (IETF) and by the users of the Internet itself in the form of Requests for Comments (RFCs).

1.1 IP

Internet Protocol (IP) is the core protocol of the TCP/IP suite.

RFC 791 states:

"The Internet Protocol is designed for use in interconnected systems of packet switched computer communication networks."

IP is not designed to provide any additional services beyond its primary function, which is to deliver a packet of bits (a "datagram") from point A to point B over any network it happens to encounter along the way.

IP doesn't, in and of itself, know anything about the information in the datagram it carries, nor does it have any provision beyond a simple checksum to ensure that the data is intact or that it has reached its destination. These functions are left to the other protocols that are part of the TCP/IP suite.

1.2 TCP

According to RFC 793, Transmission Control Protocol (TCP) is designed to be:

"a connection-oriented, end-to-end reliable protocol designed to fit into a layered hierarchy of protocols which support multi-network applications."

Connection-oriented

TCP provides for the communication of packets of bits between two points and connects them together, sending the datagram specifically from one computer or other device to another computer or device.

End-to-end

TCP packets have specific endpoints designated in the packet, and those packets are passed along the wire and ignored except by the actual endpoint of the packet and any device that needs to direct the packet.

Reliable

When a program, such as ftp, uses TCP for its protocol, the TCP/IP suite itself takes responsibility for the reliability of the communication. It provides for interprocess communication to ensure that the packets that are sent not only get to the intended destination, but get there in the order in which they were sent. If a packet is missing, the protocol will communicate the necessary information back to the sending device to ensure that the packet is re-sent.

Since TCP has to create a reliable connection between two devices or processes, there is substantially more overhead associated with each packet that is sent than with the other, less reliable protocols included in the suite.

UDP

User Datagram Protocol (UDP) is a connectionless, transaction-oriented protocol designed for sending packets with a minimum of protocol overhead, as defined by RFC 768.

It provides no guarantee that the packet was received by its intended recipient or that the information in the packet was received in the order in which it was sent.

UDP is frequently used for broadcast messages, where there is no specific intended recipient (such as BOOTP and DHCP requests); but it also can be used for applications where the sender is willing to spend extra, internal overhead to ensure reliable delivery, which actually has the result of decreasing the overall overhead of the underlying protocol.

Windows Sockets

Windows Sockets are a standard way of allowing applications programs to communicate with a TCP/IP stack without having to worry about the underlying variations in TCP/IP stack implementation.

In the past there were many different vendors for TCP/IP protocol and applications suites running on MS-DOS-based computers, and each vendor had a slightly different version of the TCP/IP protocol and application suites. This made it extremely difficult to write an application program using TCP/IP that would work with all possible TCP/IP implementations. The Windows Sockets interface was designed to get around this problem by providing a uniform set of Application Programming Interface (API) calls that would remain the same, regardless of the underlying differences in the actual implementation of TCP/IP.

Windows NT Server supports the current version of Windows Sockets (version 2.0); version 2 provides compatibility with earlier versions and provides improved functionality and support for additional features.

RFCs

The purpose of RFCs is to provide a way to communicate and to agree on the architecture and functionality of the Internet.

Key RFC's and their Topics

RFC791	Internet Protocol (IP)
RFC792	Internet Control Message Protocol (ICMP)
RFC793	Transmission Control Protocol (TCP)
RFC768	User Datagram Protocol (UDP)
RFC854	Telnet Protocol
RFC959	File Transfer Protocol (FTP)
RFC821	Simple Mail Transfer Protocol (SMTP)
RFC822	Standard for the format of ARPA Internet text messages
RFC1117	Assigned numbers
RFC991	Official ARPA Internet protocols
RFC1034	DNS - concepts and facilities
RFC1035	DNS - implementation and specification

IP Addresses

An IP address identifies a specific computer using a simple, 32-bit addressing scheme.

Four octets, separated by dots, in the form w.x.y.z, describe a combination of the network's address and the local machine's address on that network.

Networks are assigned one of three classes: A, B, or C. The classes divide the networks according to their size and complexity.

Class A Networks

A class A network has an address that begins with a number from 0 through 127, which represents the first w portion of the address and which describes the network itself; the remainder of the address is the address of the actual local device on that network. The class A address 127, however, is reserved for a specific function and isn't available for general use. This means there are 127 possible class A addresses in the world, and each one can contain more than 16 million unique network devices.

The Loopback Address

All IP addresses that begin with the network number 127 are very special. They are interpreted by your network card as a loopback address. Any packet sent to an address beginning with 127 is treated as if the destination address were the local device and the packet had arrived at its intended address. This means that packets addressed to 127.0.0.1 are treated the same as packets addressed to 127.37.90.17 - both addresses are actually that of your current machine. The same is true for all of the other 16 million addresses in the 127 class A network.

Class B Networks

A class B network has an address that begins with a number from 128 through 191, which represents the first w portion of the address and which describes the network itself. The remainder of the address is the address of the actual local device on that network. There are approximately 16 thousand class B networks, each of which can have 64 thousand unique addresses on it.

Many of the addresses in the class B address space have been broken up into smaller groups of addresses and reassigned. Large Internet service providers (ISPs), for example, use this technique to utilize the available address space more efficiently.

Class C Networks

A class C network has an address that begins with a number from 192 through 223, which represents the first w portion of the address and which describes the network itself. The remainder of the address is the address of the actual local device on that network. There are roughly 2 million class C networks, each of which can have a maximum of 254 devices on the network - enough address space for a small business network or a department network.

Class D and Class E Addresses

An IP address with a number from 224 through 239 for the first w portion of the address is known as a class D address. It is used for multicast addresses. An IP address with a number from 240 through 247 for the first w portion of the address is referred to as a class E address; this is space that is reserved for future use.

2 Subnets, Gateways, and Routers

If every computer on the Internet saw every packet as it was transmitted, the overhead would be overwhelming - machines would be swamped and the entire Internet would have come to a grinding, screeching halt years ago.

Clearly, there was a need for filtering and routing the packets in order for it to be possible to:

- C Print to a network printer without disrupting other networks and still be able to reach any IP address on the Internet.
- C Reach any IP address without having to know a whole lot about how to get there.

2.1 Subnets

A subnet is a portion of the network that operates like a separate network. It doesn't have to concern itself with what happens outside its specific portion of the network or "worry" that events in its portion will disrupt the rest of the network.

A subnet is usually a separate physical piece of "wire" that has only one point of contact through a router or a bridge with the other areas of the network.

The following special address ranges are the ranges provided for networks that will not be connecting to the Internet.

10.0.0.0 through 10.255.255.255	A Class A network.
172.16.0.0 through 172.31.255.255	16 contiguous Class B networks.
192.168.0.0 through 192.168.255.255	256 contiguous Class C networks.

2.2 Subnet Mask

To enable computers in one portion of the network to see and communicate directly with other computers in the same part of the network, but not with computers in other parts of the network, a subnet mask is used.

A subnet mask is an address, in the form w.x.y.z, that masks or blocks areas outside a specific area from sight.

You will be able to send a packet to that address only by first passing the packet to a gateway or router that knows where you are and also either where the other network is or whom to ask about where it might be.

2.3 Binary Masks

A subnet works by filtering only those portions of the IP address space that aren't masked by a number 1. A zero allows the respective portion of the address to be seen.

Any portion of an address that has a zero for the subnet mask will be seen, but any portion of the address that has the number 255 in it will not be seen. The subnet masks are implemented using binary numbers.

All of the subnet masks on a single portion of a network must be the same.

2.4 Gateways

A gateway is a physical device, usually a router but sometimes just another computer with more than one network card in it, that is physically connected to both portions of the network.

- C If your subnet mask is 255.255.255.0 and if the y portion of your IP address is 231, you won't be able to see an IP address on the network that has a y portion of 232.
- C If you want to send a packet to the computer at this unknown address, the packet must be sent through a gateway.

It takes your packet from the 231 portion and sends it to the 232 portion.

Thus, the device acts as a gatekeeper between the separate parts of the network, keeping the traffic from the 231 portion on the 232 segment and letting a packet through to the 231 segment only when it recognizes that the packet really belongs there.

2.5 Routers

A router is a physical device that connects to more than one physical segment of the network; it sends packets between those segments as required.

If it doesn't know where the packet goes, it knows where to ask for directions to the destination address, which is on another router. The router constantly updates its routing tables with information from other routers.

If your network is part of the Internet, the router has to be able to handle a huge number of possible routes between locations and decide instantaneously the best way to get packets from one point to another.

So far, this kind of traffic jam hasn't occurred very often, but it is becoming more and it is more of a problem. The current technology in routers also is reaching its limit in being able to calculate the best route from all possible routes when a key router fails.

The next generation of TCP/IP (known as IPng or IPv6) will help, as will new algorithms for how the calculations will be done.

2.6 Routing Protocols

The most commonly used TCP/IP address resolution protocols are:

ARP: Address Resolution Protocol x	Maps the IP address to the physical hardware address (MAC address) of that IP address, permitting you to send something to an IP address without having to know which physical device it belongs to.
RARP: Reverse Address Resolution Protocol	Maps the physical hardware address (MAC address) to the IP address, permitting you to determine the IP address when you know only the physical hardware address of the receiving machine.
Proxy ARP: Proxy Address Resolution Protocol	A method for implementing subnets on older versions of TCP/IP, which don't "understand" subnetting. This protocol is defined in RFC 1027.

3 DNS: Domain Name System

The Domain Name System was designed in the early 1980s, and became the official method for mapping IP addresses to names in 1984. Since that time there have been modifications to the overall structure of the DNS database and some of the ways it works, but the overall system is still remarkably similar to its original design.

3.1 Domain Name Space

The domain name space is an expression that is sometimes used to describe the tree-shaped structure of all domains - from the root ("." or "dot") domain to the lowest level domain in the structure. This hierarchical structure separates each part of the domain name space with a dot, which lets you know where you are in the domain hierarchy.

The root domains are the first level of the tree below the root. They describe the kinds of networks that are within their domain in two or three letters. There are both functional root domains and geographical root domains.

Refer to RFC 1591 for an overall description of the Domain Name System and RFC 1034 and RFC 1035 for the actual specifications of the system.

The Internet initially used a single, master file ("HOSTS.TXT") for domain names that was sent via ftp to everyone who had to be able to convert from numbers to names. This, obviously, generated enormous overhead, even when the Internet was still relatively small; with continued growth, the system would have become totally unworkable long ago.

The DNS, however, is a distributed database that is expandable and extensible so that more information can be added as needed. It allows local administration of local names, while maintaining overall network integrity and compliance with standards.

3.2 Obtaining a Domain Name

The process for getting a name is done through the InterNIC.

1. Decide what name you'd like to have. Come up with several alternates and variants on it—you'll probably need them.

The available short names are disappearing at a rapid rate.

2. Do some research on existing names to find out if the name you want is already in use; try the variants as well, until you find one that hasn't been taken.

The simplest way to do a name search is to go to the InterNIC registration home page and use "Whois" to see if someone is already using the name.

The InterNIC address is: <http://rs.internic.net/rs-internic.html>

3. Create the necessary DNS records on your DNS server, or have your Internet service provider (ISP) do it.

If you are connecting to the Internet through an ISP and if they will be the ones who maintain your DNS records, let them do the dirty work and set up everything.

4. Register the name with the InterNIC, using the same address as in step 2.
5. Pay when the InterNIC sends you the bill, or you'll lose your right to the name.

4 From Domain Names to IP Addresses

When you click on a link to "http://www.sysed.com" the web browser tries to go there, what actually happens? How does it find //www.sysed.com.

The answer is that it asks the primary DNS server listed on the TCP/IP properties page at your workstation. But how does the DNS server know this?

When a TCP/IP application wants to communicate with or connect to another location, it needs the address of that location. But it usually starts out knowing only the name, so it first has to convert the name to an IP address.

The first place it looks is in the local cache of information containing names and the resolved associated IP addresses. The chances are that the IP address hasn't changed in this short span of time.

However, consider a scenario where the user of the workstation hasn't been surfing the web in the last couple of days and the DNS server doesn't have any recent information about //www.sysed.com.

Well, if the DNS server doesn't know, it asks around to see if another server knows. It sends a special UDP packet out on the network and asks if any server recognizes the name. Ultimately, the query will get a response from one of the other DNS servers - even if the query has to go all the way to the authoritative DNS server for the root domain.

5 DHCP

One of the problems that is facing many organizations using the Internet these days is how to handle intermittently connected computers, such as laptops and remote computers, in a limited IP address space.

DHCP provides the capability of setting aside a block of IP addresses for computers that are connected only intermittently and then parcel out the addresses as they are required.

This allows a mobile user to connect to the network whenever necessary and automatically be assigned an appropriate address that is based on the location of the computer from which the user is connecting. This means you don't have to preassign an address to a particular remote location that might be connected only one day a week when a dial-up user calls in. You also save on the total number of addresses required for your network.

There's a useful side benefit to DHCP as well. A system administrator can set up the DHCP server to provide all of the necessary configuration details to the client computer without the client or its user having to know anything about the details of TCP/IP.

This can be a definite plus, even for permanently connected computers, because it leverages the skills of the relatively more skilled administrator. He or she has to set up the configuration information once on the server, and then a less skilled assistant can configure the machines on the network without having to know as much about TCP/IP.

6 WINS

The Windows Internet Name Service (WINS) is a way to map IP addresses to NetBIOS names. Although DNS provides all of the information about a computer name that most traditional TCP/IP applications will need to know, it actually provides too much information for what Windows NT needs most often - the NetBIOS name.

DNS works off the fully qualified domain name but for Windows Networking, all Windows NT cares about is the NetBIOS name, which is usually something like "SERVER1."

WINS provides the information Windows NT needs by mapping the IP address to the NetBIOS name. WINS' other advantage over DNS is that it is a dynamic protocol that can handle DHCP-assigned IP addresses.

In order to make WINS and Browsing work correctly when you are in an enterprise environment with subnets and multiple domains, there are three possible scenarios:

- C A single domain across a subnet boundary.
- C Multiple domains within a subnet boundary.
- C Multiple domains across a subnet boundary.

A source of potential confusion in a Microsoft network is the subtle distinction between the word browsing as it is commonly used to mean looking for or at the resources available and as it is used in the Microsoft sense of the Computer Browser network service.

6.1 A Single Domain Across a Subnet Boundary

In order to see the resources of the domain across a subnet boundary, with only TCP/IP as the protocol, it will be necessary to set up WINS servers on both sides of the router or else work with the LMHOSTS files.

In general, it is a good strategy to avoid LMHOSTS files because they have to be edited manually every time there is a network change. In addition, they don't work well with DHCP, where IP addresses can change.

In general, it's a good idea to put a backup domain controller on the far side of the router from your primary domain controller. The obvious choice is to make both the backup domain controller and the primary domain controller WINS servers.

6.2 Multiple Domains within a Subnet Boundary

To see and use the resources across a domain boundary when there are multiple domains within a subnet boundary, it will be necessary to set up a trust relationship between the domains.

It also will be necessary for both domains to have WINS server. The implementation can be set up and explicitly defined to have the other domains to browse.

6.3 Multiple Domains Across a Subnet Boundary

The Browsing packets won't cross a subnet boundary unless they know where they're going. To cross domain and router boundaries, it will be necessary to explicitly add the domains from the other side of the router that is to be browsed.

To achieve this, implement the following steps:

Part A: Modify the Browser Service

1. Open the Network dialog box, and click the Services tab.
2. Double-click Computer Browser.
3. Add the domain(s) that are on the other side of a router boundary, and click OK.

Part B: Set up the WINS Servers

Set up WINS servers in each of the domains and establish the necessary trust relationships.

7 IPng: IPv6

Several years ago, the IETF and others began working on solutions to the limitations of the 32-bit address space and the routing protocols in the current TCP/IP structure.

The various proposed solutions have sorted themselves into a consensus on the next generation of IP, known as IPng or IPv6 (version 6 of the Internet Protocol).

IPv6 was formally accepted by the IETF in December of 1994 and was documented in RFC 1752. It defines a 128-bit IP address space that is compatible with the current implementation of TCP/IP (version 4 or IPv4), but it provides for a greatly increased address space (3×10^{38}), as well as including additional information in the packets to provide for improved routing and handling of mobile devices.

It is important to recognize that IPng or IPv6 is evolutionary, not revolutionary. The transition can be made gradually because the protocol will be able to coexist in most situations with existing IPv4 implementations.

8 Administration Tools for TCP/IP

Windows NT Server provides tools for setting up and managing a:

C	DNS server	C	WINS server	C	DHCP server
---	------------	---	-------------	---	-------------

8.1 DNS Manager

DNS Manager provides a means for:

- C setting up and managing a DNS server to resolve domain names to IP addresses locally.
- C maintaining records for a domain.
- C administering multiple servers from the same application.

If DNS server is installed, DNS Manager can be opened from the Administrative Tools (Common) menu off the Start menu.

Adding a Server

To add a server to the list of DNS servers which can be administered:

1. Choose New Server from the DNS menu.
2. In the Add DNS Server dialog box, type the name or the IP address of the DNS server to be managed in the DNS Server text box.

The Windows NT Server DNS Administrator will attempt to connect to the server; and, if it is successful, it will display the statistics for that server as well as the types of records and zones maintained by the server.

Windows NT Server DNS Manager only works with Windows NT DNS servers; it can not be used to administer any other DNS servers you might be running.

DNS Server Functions

From the DNS menu, there are several functions from which to choose:

Pause Zone	Takes the zone offline while it is being worked on.
New Zone	Allows a zone to be added, either as a primary or secondary server.
Add Domain	Allows a new domain to be added.
New Host	Allows a new host to be added and create the A record and the PTR record, for the host.
New Record	Allows a new record to be added to any of the supported types for the host or domain selected.
Delete	Removes the selected item.
Update	Server data files increments the serial number and updates the database.
Properties	Shows detailed properties for the selected item.

Caveat

When making a change to the DNS records, make to choose Update Server Data Files.

This will increment the serial number, letting other DNS servers know that a change had been made to the information and that they need to update their information, which is now out of date.

Supported DNS Records

Windows NT Server DNS supports a variety of DNS record types.

DNS Address Types	Description
A Record	Address Record
AAAA Record	IPv6 Address Record
AFSDB Record	Andrews File System or DCE Record
CNAME Record	Alias Record
HINFO Record	Host Information Record
ISDN	ISDN Information Record
MB Record	Mailbox Name Record
MG Record	Mail Group Record
MINFO Record	Mailbox Information Record
MR Record	Mailbox Renamed Record
MX Record	Mail Exchange - "Smart Host" - Record
NS	Record Name Server Record
PTR Record	Pointer Resource Record
RP Record	Responsible Person Information Record
RT Record	Route-Through Record
TXT Record	Text Record
WKS Record	Well Known Services Record
X25 Record	X.25 Information Record

8.2 WINS Manager

WINS Manager provides the capability for setting up and managing a WINS server to convert IP addresses into NetBIOS names.

When a WINS Server is installed, WINS Manager can be opened from the Administrative Tools (Common) menu off the Start menu.

Adding a Server

A WINS server can be added to the list of WINS servers to be managed. Choose Add WINS Server from the Server menu.

In the WINS Servers list box of the WINS Manager window, by default, the primary and secondary WINS servers for the local computer will be visible.

To add a server, type the name or the IP address of the server in the WINS Server text box. In the right pane of the WINS Manager window, the current statistics for the selected WINS server will be visible.

WINS Server Functions

WINS Manager provides for the management of the WINS operations on multiple servers from the same application.

The supported functions are:

Function	Description
Add WINS Server	Add other WINS servers that you want to manage.
Delete WINS Server	Removes the selected WINS server from local management.
Detailed Information	Displays detailed statistics about the selected WINS server.
Configuration	Modifies server configuration information, including expiration times, replication parameters, and logging.
Replication Partners	Allows for the addition or deletion of replication partners and set options for the replication.
Show Database	Shows the full WINS database, including static mappings.
Initiate Scavenging	Purges the WINS database, and generally cleans it up.
Static Mappings	Allows for adding assigned names to fixed IP addresses manually.

All of these functions can be performed on multiple WINS servers, not just the one from which the application the application is running.

WINS Manager can also be run from any Windows NT version 4 client.

8.3 DHCP Manager

DHCP Manager provides for setting up and managing a DHCP server to assign and manage IP addresses and their properties for DHCP clients.

If DHCP server is installed, DHCP Manager can be opened from the Administrative Tools (Common) menu off the Start menu.

WINS Show Database dialog box showing the NetB10S-name-to-IP-Address mappings of a WINS server.

Adding a Server

A DHCP server can be added to the list of DHCP servers to be managed.

1. Choose Add from the Server menu.
2. In the DHCP Server text box, type the name or IP address of the server to be added, and then click OK.

DHCP Server

DHCP Manager provides for the administration of all the DHCP servers, and their properties and functionality, from a single location.

These functions include the following:

Function	Description
Remove	Removes the selected server from the list of managed servers.
Create Scope	Creates a new scope for the selected subnet.
Scope Properties	Edits the properties for the selected scope.
Deactivate or Delete the Scope	Temporarily deactivates or reactivates a scope on the server or permanently deletes the scope.
Add Reservations	Adds reserved addresses for particular clients to the scope.
Active Leases	Shows all active client leases on the server, including reserved addresses.
DHCP Options	<ul style="list-style-type: none"> Ⓒ Scope Edits the configuration options for the selected scope. Ⓒ Global Edits the global client configuration options. Ⓒ Default Edits the default client configuration options.

The DHCP server allows many options to be preset that would normally need to be set manually for a standard, fixed-address, TCP/IP device.

With DHCP Manager many of the options can be set individually, for each scope, or globally. Individual clients can override the default settings, of course; but in most cases this will be neither necessary nor desirable, if the DHCP options have been set correctly.

The following options can be configured on a per scope, global, or default basis:

Time Offset	Policy Filter Masks	Default TTL Option
Router	Non-Local Source Routing	Keepalive Interval
Timer Server	Maximum DG Reassembly Size	Keepalive Garbage
Name Servers	Default Time to Live	NIS Domain Name
DNS Servers	Path MTU Aging TO	NIS Servers
Log Servers	Path MTU Plateau Table	NTP Servers
Cookie Servers	MTU Options	Vendor Specific Info
LPR Servers	All Subnets Are Local (for MTU)	WINS/NBNS Servers
Impress Servers	Broadcast Address	WINS/NBT Node Type
Resource Location Servers	Perform Mask Discovery	NetBIOS Scope ID
Host Name	Mask Supplier Option	XWindow System Font
Boot File Size	Perform Router Discovery	XWindow System Display
Merit Dump File	Router Solicitation Address	NIS+ Domain Name
Swap Server	Static Route Option	NIS+ Servers
Root Path	Trailer Encapsulation	Bootfile Name
Extensions Path	ARP Cache Timeout	Mobile IP Home Agents
IP Layer Forwarding	Ethernet Encapsulation	