TCP/IP Networking Basics

NETGEAR

September 2005

Trademarks

NETGEAR and Auto Uplink are trademarks or registered trademarks of NETGEAR, Inc..

Microsoft, Windows, and Windows NT are registered trademarks of Microsoft Corporation.

Other brand and product names are registered trademarks or trademarks of their respective holders. Portions of this document are copyright Intoto, Inc.

Statement of Conditions

In the interest of improving internal design, operational function, and/or reliability, NETGEAR reserves the right to make changes to the products described in this document without notice.

NETGEAR does not assume any liability that may occur due to the use or application of the product(s) or circuit layout(s) described herein.

Product and Publication Details

Model Number:	n/a
Publication Date:	September 2005
Product Family:	Product Family
Product Name:	n/a
Home or Business Product:	n/a
Language:	English
Publication Version Number:	1.0

Contents

TCP/IP Networking Basics

Chapter 1
About This Manual

This chapter describes the intended audience, scope, conventions, and formats of this manual.

Audience, Scope, Conventions, and Formats

This manual assumes that the reader has basic to intermediate computer and Internet skills. However, basic computer network, Internet, and firewall technologies tutorial information is provided on the NETGEAR Web site.

This manual uses the following typographical conventions:

Table 1-1. Typographical Conventions

italics	Emphasis, books, CDs, URL names
bold	User input
fixed font	Screen text, file and server names, extensions, commands, IP addresses

This manual uses the following formats to highlight special messages:

> **Note:** This format is used to highlight information of importance or special interest.

> **Tip:** This format is used to highlight a procedure that will save time or resources.

How to Use this Manual

The HTML version of this manual includes the following:

- Buttons, $\boxed{\text{ > }}$ and $\boxed{\text{ < }}$, for browsing forwards or backwards through the manual one page at a time

- A $\boxed{\equiv}$ button that displays the table of contents. Double-click on a link in the table of contents to navigate directly to where the topic is described in the manual.

- A $\boxed{\text{▱}}$ button to access the full NETGEAR, Inc. online knowledge base for the product model.

- Links to PDF versions of the full manual and individual chapters.

How to Print this Manual

To print this manual you can choose one of the following several options, according to your needs.

- **Printing a Page in the HTML View**.

  Each page in the HTML version of the manual is dedicated to a major topic. Use the *Print* button on the browser toolbar to print the page contents.

- **Printing a Chapter**.

  Use the *PDF of This Chapter* link at the top left of any page.

  — Click the "*PDF of This Chapter*" link at the top right of any page in the chapter you want to print. The PDF version of the chapter you were viewing opens in a browser window.

  — Your computer must have the free Adobe Acrobat reader installed in order to view and print PDF files. The Acrobat reader is available on the Adobe Web site at *http://www.adobe.com*.

  — Click the print icon in the upper left of the window.

> **Tip:** If your printer supports printing two pages on a single sheet of paper, you can save paper and printer ink by selecting this feature.

- **Printing the Full Manual**.

    Use the *Complete PDF Manual* link at the top left of any page.

    — Click the *Complete PDF Manual* link at the top left of any page in the manual. The PDF
      version of the complete manual opens in a browser window.

    — Click the print icon in the upper left of the window.

    > **Tip:** If your printer supports printing two pages on a single sheet of paper, you can
    > save paper and printer ink by selecting this feature.

Chapter 2
TCP/IP Networking Basics

A network in your home or small business uses the same type of TCP/IP networking that is used for the Internet. This manual provides an overview of IP (Internet Protocol) networks and networking.

Related Publications

As you read this document, you may be directed to RFC documents for further information. An RFC is a Request For Comment (RFC) published by the Internet Engineering Task Force (IETF), an open organization that defines the architecture and operation of the Internet. The RFC documents explain the standard protocols and procedures for the Internet. The documents are listed on the World Wide Web at *http://www.ietf.org* and can also be found on many other Web sites.

Basic Router Concepts

Large amounts of bandwidth can be provided easily and relatively inexpensively in a local area network (LAN). However, providing high bandwidth between a local network and the Internet can be very expensive. Because of this expense, Internet access is usually provided by a slower-speed wide-area network (WAN) link such as a cable or DSL modem. For the WAN link to work on the Internet, the data traffic meant for the Internet needs to be separated from other WAN data and forwarded. A router usually performs the tasks of selecting and forwarding this data.

What is a Router?

A router is a device that forwards traffic between networks based on network layer information in the data and on routing tables maintained by the router. In these routing tables, a router builds up a logical picture of the overall network by gathering and exchanging information with other routers in the network. Using this information, the router chooses the best path for forwarding network traffic.

Routers vary in performance and scale, number of routing protocols supported, and types of physical WAN connection they support.

Routing Information Protocol

One of the protocols used by a router to build and maintain a picture of the network is the Routing Information Protocol (RIP). Using RIP, routers periodically update one another and check for changes to add to the routing table. RIP-2 supports subnet and multicast protocols. RIP is not required for most home applications.

Internet Protocol (IP) Addresses

Because TCP/IP networks are interconnected across the world, each computer on the Internet must have a unique address (called an *IP address*) to make sure that transmitted data reaches the correct destination. Blocks of addresses are assigned to organizations by the Internet Assigned Numbers Authority (IANA). Individual users and small organizations may obtain their addresses either from the IANA or from an Internet service provider (ISP). You can contact IANA at *http://www.iana.org*.

The Internet Protocol (IP) uses a 32-bit address structure. The address is usually written in dot notation (also called dotted-decimal notation), in which each group of eight bits is written in decimal form, separated by decimal points.

For example, the following binary address:

```
11000011  00100010  00001100  00000111
```

is normally written as:

```
195.34.12.7
```

The latter version is easier to remember and easier to enter into your computer.

In addition, the 32 bits of the address are subdivided into two parts. The first part of the address identifies the network, and the second part identifies the host node or station on the network. The dividing point may vary depending on the address range and the application.

There are five standard classes of IP addresses. These address classes have different ways of determining the network and host sections of the address, allowing for different numbers of hosts on a network. Each address type begins with a unique bit pattern, which is used by the TCP/IP software to identify the address class. After the address class has been determined, the software can correctly identify the host section of the address. The figure below shows the three main address classes, including network and host sections of the address for each address type.

Class A

Network   Node

Class B

Network   Node

Class C

Network   Node

7261
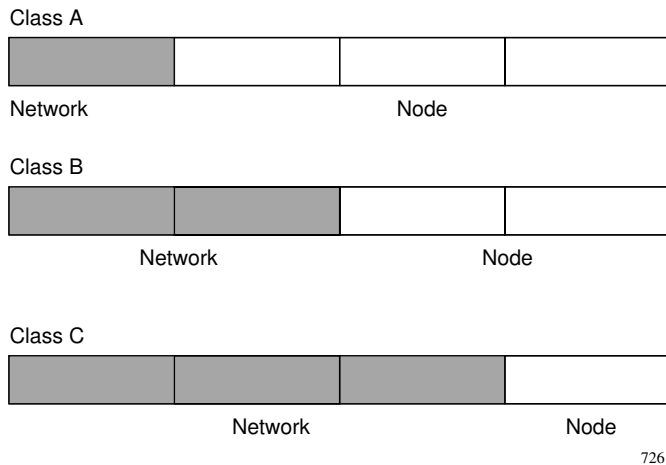
Figure 2-1

The five address classes are:

*   **Class A**
    Class A addresses can have up to 16,777,214 hosts on a single network. They use an 8-bit network number and a 24-bit node number. Class A addresses are in this range:

    ```
     1.x.x.x to 126.x.x.x.
    ```

*   **Class B**
    Class B addresses can have up to 65,354 hosts on a network. A Class B address uses a 16-bit network number and a 16-bit node number. Class B addresses are in this range:

    ```
     128.1.x.x to 191.254.x.x.
    ```

*   **Class C**
    Class C addresses can have up to 254 hosts on a network. A Class C address uses a 24-bit network number and an 8-bit node number. Class C addresses are in this range:

    ```
     192.0.1.x to 223.255.254.x.
    ```

*   **Class D**
    Class D addresses are used for multicasts (messages sent to many hosts). Class D addresses are in this range:

    ```
     224.0.0.0 to 239.255.255.255.
    ```

*   **Class E**
    Class E addresses are for experimental use.

This addressing structure allows IP addresses to uniquely identify each physical network and each node on each physical network.

v1.0, September 2005

For each unique value of the network portion of the address, the base address of the range (host address of all zeros) is known as the *network address* and is not usually assigned to a host. Also, the top address of the range (host address of all ones) is not assigned, but is used as the broadcast address for simultaneously sending a packet to all hosts with the same network address.

Netmask

In each of the address classes previously described, the size of the two parts (network address and host address) is implied by the class. This partitioning scheme can also be expressed by a netmask associated with the IP address. A netmask is a 32-bit quantity that, when logically combined (using an AND operator) with an IP address, yields the network address. For instance, the netmasks for Class A, B, and C addresses are 255.0.0.0, 255.255.0.0, and 255.255.255.0, respectively.

For example, the address 192.168.170.237 is a Class C IP address whose network portion is the upper 24 bits. When combined (using an AND operator) with the Class C netmask, as shown here, only the network portion of the address remains:

```
11000000   10101000   10101010   11101101   (192.168.170.237)
```

combined with:

```
11111111   11111111   11111111   00000000   (255.255.255.0)
```

equals:

```
11000000   10101000   10101010   00000000   (192.168.170.0)
```

As a shorter alternative to dotted-decimal notation, the netmask may also be expressed in terms of the number of ones from the left. This number is appended to the IP address, following a backward slash (/), as "/n." In the example, the address could be written as 192.168.170.237/24, indicating that the netmask is 24 ones followed by 8 zeros.

Subnet Addressing

By looking at the addressing structures, you can see that even with a Class C address, there are a large number of hosts per network. Such a structure is an inefficient use of addresses if each end of a routed link requires a different network number. It is unlikely that the smaller office LANs would have that many devices. You can resolve this problem by using a technique known as subnet addressing.

Subnet addressing allows us to split one IP network address into smaller multiple physical networks known as subnetworks. Some of the node numbers are used as a subnet number instead. A Class B address gives us 16 bits of node numbers translating to 64,000 nodes. Most organizations do not use 64,000 nodes, so there are free bits that can be reassigned. Subnet addressing makes use of those bits that are free, as shown below.

Class B

7262
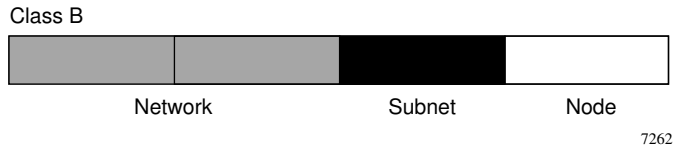
Figure 2-2

A Class B address can be effectively translated into multiple Class C addresses. For example, the IP address of 172.16.0.0 is assigned, but node addresses are limited to 255 maximum, allowing eight extra bits to use as a subnet address. The IP address of 172.16.97.235 would be interpreted as IP network address 172.16, subnet number 97, and node number 235. In addition to extending the number of addresses available, subnet addressing provides other benefits. Subnet addressing allows a network manager to construct an address scheme for the network by using different subnets for other geographical locations in the network or for other departments in the organization.

Although the preceding example uses the entire third octet for a subnet address, note that you are not restricted to octet boundaries in subnetting. To create more network numbers, you need only shift some bits from the host address to the network address. For instance, to partition a Class C network number (192.68.135.0) into two, you shift one bit from the host address to the network address. The new netmask (or subnet mask) is 255.255.255.128. The first subnet has network number 192.68.135.0 with hosts 192.68.135.1 to 129.68.135.126, and the second subnet has network number 192.68.135.128 with hosts 192.68.135.129 to 192.68.135.254.

> **Note:** The number 192.68.135.127 is not assigned because it is the broadcast address of the first subnet. The number 192.68.135.128 is not assigned because it is the network address of the second subnet.

The following table lists the additional subnet mask bits in dotted-decimal notation. To use the table, write down the original class netmask and replace the 0-value octets with the dotted-decimal value of the additional subnet bits. For example, to partition your Class C network with subnet mask 255.255.255.0 into 16 subnets (four bits), the new subnet mask becomes 255.255.255.240.

Table 2-1. Netmask Notation Translation Table for One Octet

Number of Bits	Dotted-Decimal Value
1	128
2	192
3	224
4	240
5	248
6	252
7	254
8	255

The following table displays several common netmask values in both the dotted-decimal and the masklength formats.

Table 2-2. Netmask Formats

Dotted-Decimal	Masklength
255.0.0.0	/8
255.255.0.0	/16
255.255.255.0	/24
255.255.255.128	/25
255.255.255.192	/26
255.255.255.224	/27
255.255.255.240	/28
255.255.255.248	/29
255.255.255.252	/30
255.255.255.254	/31
255.255.255.255	/32

Configure all hosts on a LAN segment to use the same netmask for the following reasons:

- So that hosts recognize local IP broadcast packets

  When a device broadcasts to its segment neighbors, it uses a destination address of the local network address with all ones for the host address. In order for this scheme to work, all devices on the segment must agree on which bits comprise the host address.

- So that a local router or bridge recognizes which addresses are local and which are remote

Private IP Addresses

If your local network is isolated from the Internet (for example, when using Network Address Translation, NAT, which is described below), you can assign any IP addresses to the hosts without problems. However, the IANA has reserved the following three blocks of IP addresses specifically for private networks:

```
10.0.0.0 - 10.255.255.255
172.16.0.0 - 172.31.255.255
192.168.0.0 - 192.168.255.255
```

Choose your private network number from this range. Some NETGEAR products have DHCP servers that are preconfigured to automatically assign private addresses.

Regardless of your particular situation, do not create an arbitrary IP address; always follow the guidelines explained here. For more information about address assignment, refer to RFC 1597, *Address Allocation for Private Internets,* and RFC 1466, *Guidelines for Management of IP Address Space*. The Internet Engineering Task Force (IETF) publishes RFCs on its Web site at *http://www.ietf.org*.

Single IP Address Operation Using NAT

In the past, if multiple computers on a LAN needed to access the Internet simultaneously, you had to obtain a range of IP addresses from the ISP. This type of Internet account is more costly than a single-address account typically used by a single user with a modem, rather than a router. NETGEAR products use an address-sharing method called Network Address Translation (NAT). This method allows several networked computers to share an Internet account using only a single IP address, which may be statically or dynamically assigned by your ISP.

The router does this by translating the internal LAN IP addresses to a single address that is unique on the Internet. The internal LAN IP addresses can be either private addresses or registered addresses. For more information about IP address translation, refer to RFC 1631, *The IP Network Address Translator (NAT)*.

The following figure illustrates a single IP address operation.

Private IP addresses
assigned by user

IP addresses
assigned by ISP

192.168.0.2

192.168.0.3

192.168.0.1   172.21.15.105

Internet

192.168.0.4

192.168.0.5
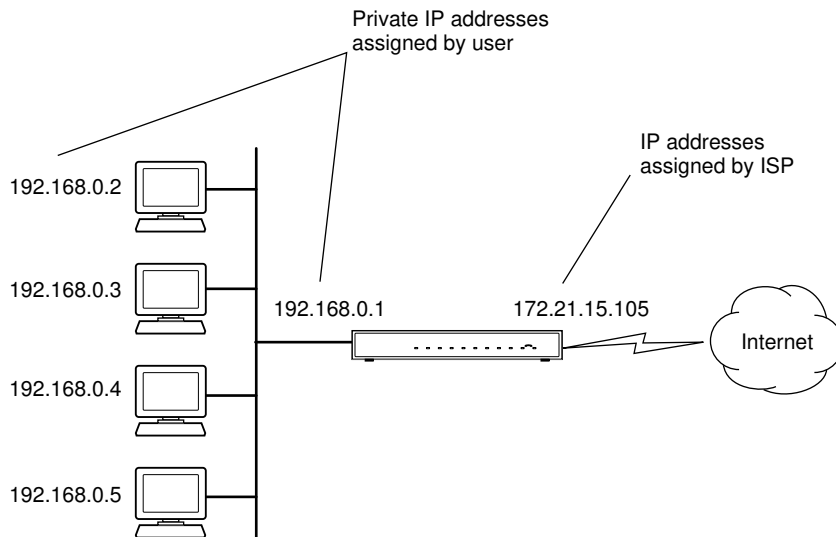
7786EA

Figure 2-3

This scheme offers the additional benefit of firewall-like protection because the internal LAN addresses are not shown to the Internet connection. This filtering can prevent intruders from probing your system. However, using port forwarding, you can allow one computer (for example, a Web server) on your local network to be accessible to outside users.

Media Access Control (MAC) Addresses and Address Resolution Protocol

An IP address alone cannot be used to deliver data from one LAN device to another. To send data between LAN devices, you must convert the IP address of the destination device to its MAC address. Each device on an Ethernet network has a unique MAC address, which is a 48-bit number assigned to each device by the manufacturer. The technique that associates the IP address with a MAC address is known as *address resolution*. Internet Protocol uses the Address Resolution Protocol (ARP) to resolve MAC addresses.

If a device sends data to another station on the network and the destination MAC address is not yet recorded, ARP is used. An ARP request is broadcast onto the network. All stations (computers, for example) on the network receive and read the request. The destination IP address for the chosen station is included as part of the message so that only the station with this IP address responds to the ARP request. All other stations discard the request.

Related Documents

The station with the correct IP address responds with its own MAC address directly to the sending device. The receiving station provides the transmitting station with the required destination MAC address. The IP address data and MAC address data for each station are held in an ARP table. The next time data is sent, the address can be obtained from the address information in the table.

For more information about address assignment, refer to the IETF documents RFC 1597, *Address Allocation for Private Internets,* and RFC 1466, *Guidelines for Management of IP Address Space*.

For more information about IP address translation, refer to RFC 1631, *The IP Network Address Translator (NAT)*.

Domain Name System (DNS) Server

Many of the resources on the Internet can be addressed by simple descriptive names such as *http://www.NETGEAR.com*. This addressing is very helpful at the application level, but the descriptive name must be translated to an IP address in order for a user to actually contact the resource. Just as a telephone directory maps names to phone numbers, or as an ARP table maps IP addresses to MAC addresses, a DNS server maps descriptive names of network resources to IP addresses.

When a computer accesses a resource by its descriptive name, it first contacts a DNS server to obtain the IP address of the resource. The computer sends the desired message using the IP address. Many large organizations, such as ISPs, maintain their own DNS servers and allow their customers to use the servers to look up addresses.

IP Configuration by DHCP

When an IP-based local area network is installed, each computer must be configured with an IP address. If the computers need to access the Internet, they should also be configured with a gateway address and one or more DNS server addresses. As an alternative to manual configuration, Dynamic Host Configuration Protocol (DHCP) is a method by which each computer on the network can automatically obtain this configuration information. A device on the network may act as a DHCP server. The DHCP server stores a list or pool of IP addresses, along with other information (such as gateway and DNS addresses) that it may assign to the other devices on the network. Some NETGEAR products can act as DHCP servers.

Some NETGEAR products also function as DHCP clients when connecting to the ISP. Such NETGEAR products can automatically obtain an IP address, subnet mask, DNS server addresses, and a gateway address if the ISP provides this information by DHCP.

Internet Security and Firewalls

When your LAN connects to the Internet through a router, an opportunity is created for outsiders to access or disrupt your network. A NAT router provides some protection because by the very nature of the process, the network behind the router is shielded from access by outsiders on the Internet. However, there are methods by which a determined hacker can possibly obtain information about your network or at the least can disrupt your Internet access. A greater degree of protection is provided by a firewall router.

What is a Firewall?

A firewall is a device that protects one network from another while allowing communication between the two. A firewall incorporates the functions of the NAT router, while adding features for dealing with a hacker intrusion or attack. Several known types of intrusion or attack can be recognized when they occur. When an incident is detected, the firewall can log details of the attempt, and it can optionally send e-mail to an administrator to report the incident. Using information from the log, the administrator can take action with the ISP of the hacker. In some types of intrusions, the firewall can fend off the hacker by discarding all further packets from the hacker's IP address for a period of time.

Stateful Packet Inspection

Unlike simple Internet sharing routers, a firewall uses a process called *stateful packet inspection* to ensure secure firewall filtering to protect your network from attacks and intrusions. Since user-level applications such as FTP and Web browsers can create complex patterns of network traffic, it is necessary for the firewall to analyze groups of network connection states. Using stateful packet inspection, an incoming packet is intercepted at the network layer and then analyzed for state-related information associated with all network connections. A central cache within the firewall keeps track of the state information associated with all network connections. All traffic passing through the firewall is analyzed against the state of these connections to determine whether or not it will be allowed to pass through or be rejected.

Denial of Service Attack

A hacker may be able to prevent your network from operating or communicating by launching a Denial of Service (DoS) attack. The method used for such an attack can be as simple as merely flooding your site with more requests than it can handle. A more sophisticated attack may attempt to exploit some weakness in the operating system used by your router or gateway. Some operating systems can be disrupted by simply sending a packet with incorrect length information.

Ethernet Cabling

Although Ethernet networks originally used thick or thin coaxial cable, most installations currently use unshielded twisted pair (UTP) cabling. The UTP cable contains eight conductors, arranged in four twisted pairs, and is terminated with an RJ45 type connector. A normal straight-through UTP Ethernet cable follows the EIA568B standard wiring as described below.

Table 2-1. UTP Ethernet cable wiring, straight-through

Pin	Wire color	Signal
1	Orange/White	Transmit (Tx) +
2	Orange	Transmit (Tx) -
3	Green/White	Receive (Rx) +
4	Blue	
5	Blue/White	
6	Green	Receive (Rx) -
7	Brown/White	
8	Brown	

Category 5 Cable Quality

Category 5 distributed cable that meets ANSI/EIA/TIA-568-A building wiring standards can be a maximum of 328 feet (ft.) or 100 meters (m) in length, divided as follows:

20 ft. (6 m) between the hub and the patch panel (if used)

295 ft. (90 m) from the wiring closet to the wall outlet

10 ft. (3 m) from the wall outlet to the desktop device

The patch panel and other connecting hardware must meet the requirements for 100-Mbps operation (Category 5). Only 0.5 inch (1.5 cm) of untwist in the wire pair is allowed at any termination point.

A twisted pair Ethernet network operating at 10 Mbits/second (10BASE-T) will often tolerate low-quality cables, but at 100 Mbits/second (10BASE-Tx) the cable must be rated as Category 5, or Cat 5, by the Electronic Industry Association (EIA). This rating will be printed on the cable jacket. A Category 5 cable will meet specified requirements regarding loss and crosstalk. In addition, there are restrictions on maximum cable length for both 10- and 100-Mbits/second networks.

Inside Twisted Pair Cables

For two devices to communicate, the transmitter of each device must be connected to the receiver of the other device. The crossover function is usually implemented internally as part of the circuitry in the device. Computers and workstation adapter cards are usually media-dependent interface ports, called MDI or uplink ports. Most repeaters and switch ports are configured as media-dependent interfaces with built-in crossover ports, called MDI-X or normal ports. Auto Uplink™ technology automatically senses which connection, MDI or MDI-X, is needed and makes the right connection.

The figure below illustrates straight-through twisted pair cable.
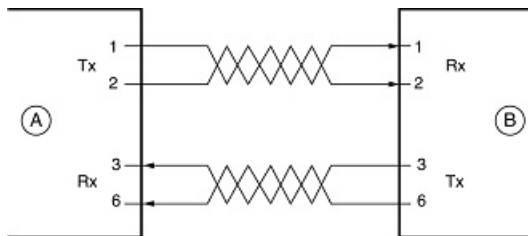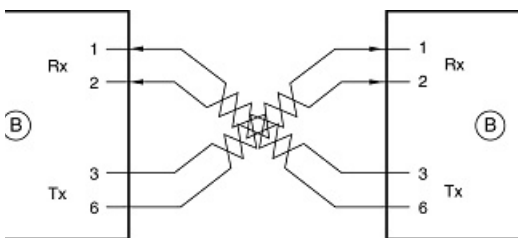

Key:
A = UPLINK OR MDI PORT (as on a PC)
B = Normal or MDI-X port (as on a hub or switch)
1, 2, 3, 6 = Pin numbers

Figure 2-4

The figure below illustrates crossover twisted pair cable.

Key:
B = Normal or MDI-X port (as on a hub or switch)
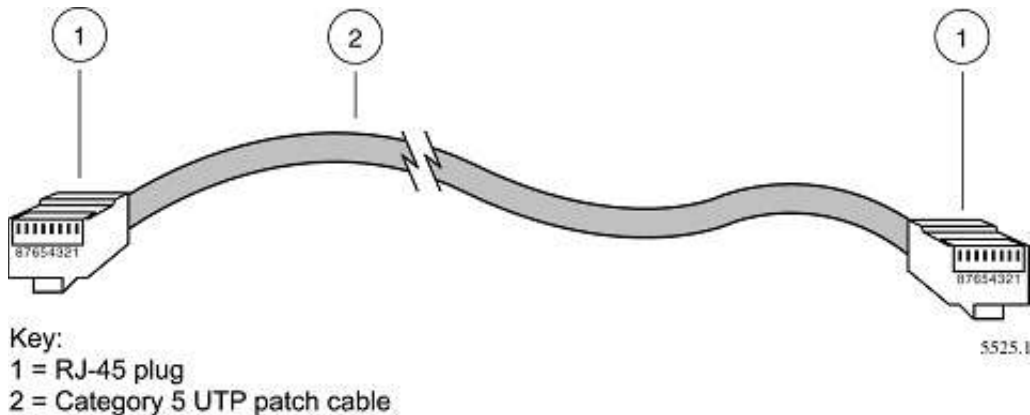1, 2, 3, 6 = Pin numbers

Figure 2-5

Key:
1 = RJ-45 plug
2 = Category 5 UTP patch cable

Figure 2-6

→ **Note:** Flat "silver satin" telephone cable may have the same RJ-45 plug. However, using telephone cable results in excessive collisions, causing the attached port to be partitioned or disconnected from the network.

Uplink Switches, Crossover Cables, and MDI/MDIX Switching

In Table 2-1 above on page 2-12, the concept of transmit and receive are from the perspective of the computer, which is wired as Media Dependant Interface (MDI). In this wiring, the computer transmits on pins 1 and 2. At the hub, the perspective is reversed, and the hub receives on pins 1 and 2. This wiring is referred to as Media Dependant Interface - Crossover (MDI-X).

When connecting a computer to a computer, or a hub port to another hub port, the transmit pair must be exchanged with the receive pair. This exchange is done by one of two mechanisms. Most hubs provide an Uplink switch which will exchange the pairs on one port, allowing that port to be connected to another hub using a normal Ethernet cable.

The second method is to use a crossover cable, which is a special cable in which the transmit and receive pairs are exchanged at one of the two cable connectors. Crossover cables are often unmarked as such, and must be identified by comparing the two connectors. Since the cable connectors are clear plastic, it is easy to place them side by side and view the order of the wire colors on each. On a straight-through cable, the color order will be the same on both connectors. On a crossover cable, the orange and blue pairs will be exchanged from one connector to the other.

Some NETGEAR products incorporate Auto Uplink technology (also called MDI/MDIX). With this feature, each *local* Ethernet port automatically senses whether the Ethernet cable plugged into the port should have a normal connection (for example, connecting to a computer) or an uplink connection (for example, connecting to a router, switch, or hub). That port then configures itself to the correct configuration. This feature also eliminates the need to worry about crossover cables because Auto Uplink will accommodate either type of cable to make the right connection.