How Routers Works

by Curt Franklin

The Internet is one of the 20th century's greatest communications developments. It allows people around the world to send e-mail to one another in a matter of seconds, and it lets you read, among other things, the articles on HowStuffWorks.com. We're all used to seeing the various parts of the Internet that come into our homes and offices -- the Web pages, e-mail messages and downloaded files that make the Internet a dynamic and valuable medium. But none of these parts would ever make it to your computer without a piece of the Internet that you've probably never seen. In fact, most people have never stood "face to machine" with the technology most responsible for allowing the Internet to exist at all: the **router**.

Photo courtesy Newstream.com

Fujitsu GeoStream R980 industrial strength router

Routers are specialized computers that send your messages and those of every other Internet user speeding to their destinations along thousands of pathways. In this article, we'll look at how these behind-the-scenes machines make the Internet work.

Keeping the Messages Moving

When you send e-mail to a friend on the other side of the country, how does the message know to end up on your friend's computer, rather than on one of the millions of other computers in the world? Much of the work to get a message from one computer to another is done by routers, because they're the crucial devices that let messages flow **between networks**, rather than within networks.

Let's look at what a very simple router might do. Imagine a small company that makes animated 3-D graphics for local television stations. There are 10 employees of the company, each with a computer. Four of the employees are animators, while the rest are in sales, accounting and management. The animators will need to send lots of very large files back and forth to one another as they work on projects. To do this, they'll use a **network**.

When one animator sends a file to another, the very large file will use up most of the network's capacity, making the network run very slowly for other users. One of the reasons that a single intensive user can affect the entire network stems from the way that Ethernet works. Each information packet sent from a computer is seen by all the other computers on the local network. Each computer then examines the packet and decides whether it was meant for its address. This keeps the basic plan of the network simple, but has performance consequences as the size of the network or level of network activity increases. To keep the animators' work from interfering with that of the folks in the front office, the company sets up two separate networks, one for the animators and one for the rest of the company. A router links the two networks and connects both networks to the Internet.

Directing Traffic

The router is the only device that sees every message sent by any computer on either of the company's networks. When the animator in our example sends a huge file to another animator, the router looks at the recipient's address and keeps the traffic on the animator's network. When an animator, on the other hand, sends a message to the bookkeeper asking about an expense-account check, then the router sees the recipient's address and forwards the message between the two networks.

One of the tools a router uses to decide where a packet should go is a **configuration table**. A configuration table is a collection of information, including:

- Information on which connections lead to particular groups of addresses
- Priorities for connections to be used
- Rules for handling both routine and special cases of traffic

A configuration table can be as simple as a half-dozen lines in the smallest routers, but can grow to massive size and complexity in the very large routers that handle the bulk of Internet messages.

A router, then, has two separate but related jobs:

- The router ensures that information doesn't go where it's not needed. This is crucial for keeping large volumes of data from clogging the connections of "innocent bystanders."
- The router makes sure that information does make it to the intended destination.

In performing these two jobs, a router is extremely useful in dealing with two separate computer networks. It joins the two networks, passing information from one to the other and, in some cases, performing translations of various **protocols** between the two networks. It also protects the networks from one another, preventing the traffic on one

from unnecessarily spilling over to the other. As the number of networks attached to one another grows, the configuration table for handling traffic among them grows, and the processing power of the router is increased. Regardless of how many networks are attached, though, the basic operation and function of the router remains the same. Since the Internet is one huge network made up of tens of thousands of smaller networks, its use of routers is an absolute necessity.

Transmitting Packets

When you make a telephone call to someone on the other side of the country, the telephone system establishes a stable circuit between your telephone and the telephone you're calling. The circuit might involve a half dozen or more steps through copper cables, switches, fiber optics, microwaves and satellites, but those steps are established and remain constant for the duration of the call. This circuit approach means that the quality of the line between you and the person you're calling is consistent throughout the call, but a problem with any portion of the circuit -- maybe a tree falls across one of the lines used, or there's a power problem with a switch -- brings your call to an early and abrupt end. When you send an e-mail message with an attachment to the other side of the country, a very different process is used.

Internet data, whether in the form of a Web page, a downloaded file or an e-mail message, travels over a system known as a **packet-switching network**. In this system, the data in a message or file is broken up into packages about 1,500 bytes long. Each of these packages gets a wrapper that includes information on the sender's address, the receiver's address, the package's place in the entire message, and how the receiving computer can be sure that the package arrived intact. Each data package, called a **packet**, is then sent off to its destination via the best available route -- a route that might be taken by all the other packets in the message or by none of the other packets in the message. This might seem very complicated compared to the circuit approach used by the telephone system, but in a network designed for data there are two huge advantages to the packet-switching plan.

- The network can balance the load across various pieces of equipment on a millisecond-by-millisecond basis.
- If there is a problem with one piece of equipment in the network while a message is being transferred, packets can be routed around the problem, ensuring the delivery of the entire message.

The Path of a Packet

The routers that make up the main part of the Internet can **reconfigure the paths** that packets take because they look at the information surrounding the data packet, and they tell each other about line conditions, such as delays in receiving and sending data and traffic on various pieces of the network. Not all routers do so many jobs, however. Routers come in different sizes. For example:

- If you have enabled Internet connection sharing between two Windows 98-based computers, you're using one of the computers (the computer with the Internet connection) as a **simple router**. In this instance, the router does so little -- simply looking at data to see whether it's intended for one computer or the other -- that it can operate in the background of the system without significantly affecting the

other programs you might be running.

- Slightly larger routers, the sort used to connect a small office network to the Internet, will do a bit more. These routers frequently enforce rules concerning security for the office network (trying to secure the network from certain attacks). They handle enough traffic that they're generally stand-alone devices rather than software running on a server.
- The largest routers, those used to handle data at the major traffic points on the Internet, handle millions of data packets every second and work to configure the network most efficiently. These routers are large stand-alone systems that have far more in common with supercomputers than with your office server.

Routing Packets: An Example

Let's take a look at a medium-sized router -- the router we use in the HowStuffWorks office. In our case, the router only has two networks to worry about: The office network, with about 50 computers and devices, and the Internet. The office network connects to the router through an Ethernet connection, specifically a 100 base-T connection (100 base-T means that the connection is 100 megabits per second, and uses a twisted-pair cable like an 8-wire version of the cable that connects your telephone to the wall jack). There are two connections between the router and our ISP (Internet service provider). One is a T-1 connection that supports 1.5 megabits per second. The other is an ISDN line that supports 128 kilobits per second. The configuration table in the router tells it that all out-bound packets are to use the T-1 line, unless it's unavailable for some reason (perhaps a backhoe digs up the cable). If it can't be used, then outbound traffic goes on the ISDN line. This way, the ISDN line is held as "insurance" against a problem with the faster T-1 connection, and no action by a staff member is required to make the switch in case of trouble. The router's configuration table knows what to do.

In addition to routing packets from one point to another, the HowStuffWorks router has rules limiting how computers from outside the network can connect to computers inside the network, how the HowStuffWorks network appears to the outside world, and other **security** functions. While most companies also have a special piece of hardware or software called a firewall to enforce security, the rules in a router's configuration table are important to keeping a company's (or family's) network secure.

One of the crucial tasks for any router is knowing when a packet of information stays on its local network. For this, it uses a mechanism called a **subnet mask.** The subnet mask looks like an IP address and usually reads "255.255.255.0." This tells the router that all messages with the sender and receiver having an address sharing the first three groups of numbers are on the same network, and shouldn't be sent out to another network. Here's an example: The computer at address 15.57.31.40 sends a request to the computer at 15.57.31.52. The router, which sees all the packets, matches the first three groups in the address of both sender and receiver (15.57.31), and keeps the packet on the local network. (You'll learn more about how the addresses work in the next section.)

Between the time these words left the Howstuffworks.com server and the time they showed up on your monitor, they passed through several routers (it's impossible to know ahead of time exactly how many "several" might be) that helped them along the way. It's very similar to the process that gets a postal letter from your mailbox to the mailbox of a friend, with routers taking the place of the mail sorters and handlers along the way.

Knowing Where to Send Data

Routers are one of several types of devices that make up the "plumbing" of a computer network. Hubs, switches and routers all take signals from computers or networks and pass them along to other computers and networks, but a router is the only one of these devices that examines each bundle of data as it passes and makes a decision about exactly where it should go. To make these decisions, routers must first know about two kinds of information: **addresses** and **network structure**.

When a friend mails a birthday card to be delivered to you at your house, he probably uses an address that looks something like this:

Joe Smith
123 Maple Street
Smalltown, FL 45678

The address has several pieces, each of which helps the people in the postal service move the letter along to your house. The ZIP code can speed the process up; but even without the ZIP code, the card will get to your house as long as your friend includes your state, city and street address. You can think of this address as a **logical address** because it describes a way someone can get a message to you. This logical address is connected to a **physical address** that you generally only see when you're buying or selling a piece of property. The survey plot of the land and house, with latitude, longitude or section bearings, gives the legal description, or address, of the property.

Logical Addresses

Every piece of equipment that connects to a network, whether an office network or the Internet, has a physical address. This is an address that's unique to the piece of equipment that's actually attached to the network cable. For example, if your desktop computer has a network interface card (NIC) in it, the NIC has a physical address permanently stored in a special memory location. This physical address, which is also called the MAC address (for Media Access Control) has two parts, each 3 bytes long. The first 3 bytes identify the company that made the NIC. The second 3 bytes are the serial number of the NIC itself.

The interesting thing is that your computer can have several logical addresses at the same time. Of course, you're used to having several "logical addresses" bring messages to one physical address. Your mailing address, telephone number (or numbers) and home e-mail address all work to bring messages to you when you're in your house. They are simply used for different types of messages -- different networks, so to speak.

Logical addresses for computer networks work in exactly the same way. You may be using the addressing schemes, or protocols, from several different types of networks simultaneously. If you're connected to the Internet (and if you're reading this, you probably are), then you have an address that's part of the TCP/IP network protocol. If you also have a small network set up to exchange files between several family computers, then you may also be using the Microsoft NetBEUI protocol. If you connect to your company's network from home, then your computer may have an address that follows Novell's IPX/SPX protocol. All of these can coexist on your computer. Since the driver software that allows your computer to communicate with each network uses resources like memory

and CPU time, you don't want to load protocols you won't need, but there's no problem with having all the protocols your work requires running at the same time.

On the next page, you'll find out how the process works for Mac machines.

MAC Addresses

The chances are very good that you'll never see the MAC address for any of your equipment because the software that helps your computer communicate with a network takes care of matching the MAC address to a logical address. The logical address is what the network uses to pass information along to your computer.

If you'd like to see the MAC address and logical address used by the Internet Protocol (IP) for your Windows computer, you can run a small program that Microsoft provides. Go to the "Start" menu, click on "Run," and in the window that appears, type WINIPCFG (IPCONFIG for Windows 2000/XP). When the gray window appears, click on "More Info" and you'll get this sort of information:

Windows 98 IP Configuration:

Host Name: NAMEHOWSTUFFWORKS
DNS Servers: 208.153.64.20
 208.153.0.5
Node Type: Broadcast
NetBIOS Scope ID:
IP Routing Enabled: Yes
WINS Proxy Enabled: No
NetBIOS Resolution Uses DNS: No

Ethernet adapter:

Description: PPP Adapter
Physical Address: 44-45-53-54-12-34
DHCP Enabled: Yes
IP Address: 227.78.86.288
Subnet Mask: 255.255.255.0
Default Gateway: 227.78.86.288
DHCP Server: 255.255.255.255
Primary WINS Server:
Secondary WINS Server: Lease Obtained: 01 01 80 12:00:00 AM
Lease Expires: 01 01 80 12:00:00 AM

There's a lot of information here that will vary depending on exactly how your connection to the Internet is established, but the physical address is the MAC address of the adapter queried by the program. The IP address is the logical address assigned to your connection by your ISP or network administrator. You'll see the addresses of other servers, including the DNS servers that keep track of all the names of Internet sites (so you can type "www.howstuffworks.com" rather than "216.27.61.189") and the gateway server that you connect to in order to reach the Internet. When you've finished looking at the information, click OK. (**Note:** For security reasons, some of the information about this connection to the Internet has been changed. You should be very careful about giving your computer's

information to other people -- with your address and the right tools, an unscrupulous person could, in some circumstances, gain access to your personal information and control your system. See this Question of the Day to learn more.)

Understanding the Protocols

The first and most basic job of the router is to know where to send information addressed to your computer. Just as the mail handler on the other side of the country knows enough to keep a birthday card coming toward you without knowing where your house is, most of the routers that forward an e-mail message to you don't know your computer's MAC address, but they know enough to keep the message flowing.

Routers are programmed to understand the most common network protocols. That means they know the format of the addresses, how many bytes are in the basic package of data sent out over the network, and how to make sure all the packages reach their destination and get reassembled. For the routers that are part of the Internet's main "backbone," this means looking at, and moving on, millions of information packages every second. And simply moving the package along to its destination isn't all that a router will do. It's just as important, in today's computerized world, that they keep the message flowing by the **best possible route**.

In a modern network, every e-mail message is broken up into small pieces. The pieces are sent individually and reassembled when they're received at their final destination. Because the individual pieces of information are called packets and each packet can be sent along a different path, like a train going through a set of switches, this kind of network is called a **packet-switched network**. It means that you don't have to build a dedicated network between you and your friend on the other side of the country. Your e-mail flows over any one of thousands of different routes to get from one computer to the other.

Depending on the time of day and day of the week, some parts of the huge public packet-switched network may be busier than others. When this happens, the routers that make up this system will communicate with one another so that traffic not bound for the crowded area can be sent by less congested network routes. This lets the network function at full capacity without excessively burdening already-busy areas. You can see, though, how **Denial of Service attacks** (described in the next section), in which people send millions and millions of messages to a particular server, will affect that server and the routers forwarding message to it. As the messages pile up and pieces of the network become congested, more and more routers send out the message that they're busy, and the entire network with all its users can be affected.

Tracing a Message

If you're using a Microsoft Windows-based system, you can see just how many routers are involved in your Internet traffic by using a program you have on your computer. The program is called **Traceroute**, and that describes what it does -- it traces the route that a packet of information takes to get from your computer to another computer connected to the Internet. To run this program, click on the "MS-DOS Prompt" icon on the "Start" menu. Then, at the **"C:\WINDOWS>"** prompt, type **"tracert www.howstuffworks.com"**. When I did this from my office in Florida, the results looked like this:

```
Tracing route to howstuffworks.com [216.27.61.189]
over a maximum of 30 hops:

 1  129 ms  133 ms  133 ms  209.215.248.11
 2  135 ms  132 ms  138 ms  209.215.248.62
 3  128 ms  112 ms  132 ms  205.152.46.184
 4  133 ms  126 ms  144 ms  205.152.48.89
 5  119 ms  112 ms  126 ms  205.152.111.248
 6  122 ms  410 ms  397 ms  903.Hssi4-0-0.GW1.ORL1.ALTER.NET [157.130.66.53]

 7  137 ms  139 ms  145 ms  115.ATM4-0.XR1.ATL1.ALTER.NET [146.188.232.194]

 8  192 ms  132 ms  138 ms  195.ATM9-0-0.BR1.ATL1.ALTER.NET [146.188.232.57]

 9  188 ms  145 ms  166 ms  sl-bb1-atl-12-0-0.sprintlink.net [144.232.9.177]

10  209 ms  171 ms  211 ms  sl-bb10-atl-0-2.sprintlink.net [144.232.12.1]
11  258 ms  158 ms  178 ms  sl-bb11-rly-5-0.sprintlink.net [144.232.9.197]
12  249 ms  179 ms  172 ms  sl-gw8-rly-4-0-0.sprintlink.net [144.232.7.246]

13  216 ms  218 ms  186 ms  sl-interlan-1-0-0.sprintlink.net [144.232.190.130]

14  232 ms  172 ms  178 ms  216.27.0.13
15  184 ms  185 ms  186 ms  howstuffworks.com [216.27.61.189]

Trace complete.
```

The first number shows how many routers are between your computer and the router shown. In this instance, there were a total of 14 routers involved in the process (number 15 is the Howstuffworks.com Web server). The next three numbers show how long it takes a packet of information to move from your computer to the router shown and back again. Next, in this example, starting with step six, comes the "name" of the router or server. This is something that helps people looking at the list but is of no importance to the routers and computers as they move traffic along the Internet. Finally, you see the Internet Protocol (IP) address of each computer or router. The final picture of this trace route shows that there were 14 routers between the Web server and me and that it took, on average, a little more than 2.5 seconds for information to get from my computer to the server and back again.

You can use Traceroute to see how many routers are between you and any other computer you can name or know the IP address for. It can be interesting to see how many steps are required to get to computers outside your nation. Since I live in the United States, I decided to see how many routers were between my computer and the Web server for the British Broadcasting Corporation. At the **C:\WINDOWS>** prompt, I typed **tracert www.bbc.com**. The result was this:

```
Tracing route to www.bbc.net.uk [212.58.240.32]
over a maximum of 30 hops:

  1   133 ms  130 ms  130 ms  209.215.248.11
  2   134 ms  110 ms  130 ms  209.215.248.61
  3   118 ms  123 ms  117 ms  205.152.46.248
  4   598 ms  760 ms  782 ms  205.152.48.89
  5  1139 ms 1244 ms  226 ms  205.152.111.248
  6   153 ms  130 ms  156 ms  Serial4-1-0.GW1.ORL1.ALTER.NET [157.130.65.157]

  7   204 ms  137 ms  148 ms  115.ATM4-0.XR2.ATL1.ALTER.NET [146.188.232.206]

  8   202 ms  170 ms  239 ms  194.ATM2-0.TR2.ATL1.ALTER.NET [146.188.232.98]
  9   150 ms  157 ms  150 ms  109.ATM6-0.TR2.DCA8.ALTER.NET [146.188.138.190]

 10   169 ms  158 ms  178 ms  296.ATM7-0.XR2.TCO1.ALTER.NET [152.63.32.217]
 11   146 ms  171 ms  158 ms  192.ATM9-0-0.GW2.TCO1.ALTER.NET [146.188.160.61]

 12   213 ms  151 ms  200 ms  extensibility-gw.customer.alter.net [157.130.4.130]
 13   236 ms  170 ms  157 ms  lo0.mp1.Washington1.level3.net [209.247.8.249]
 14   194 ms  176 ms  174 ms  loopback0.hsipaccess1.NewYork1.Level3.net [209.2.44.2.210]

 15   220 ms  157 ms  164 ms  209.244.160.70
 16   379 ms  371 ms  303 ms  www.bbc.net.uk [212.58.240.32]

Trace complete.
```

You can see that it took only one more step to reach a Web server on the other side of the Atlantic Ocean than it did to reach a server two states away!

On the next page, we'll go into detail about Denial of Service attacks.

Denial of Service Attacks

In the first quarter of 2000, there were several attacks on very popular Web sites. Most of these were "Denial of Service" attacks -- attacks that served to prevent regular readers and customers of the sites from getting a response to their requests. How did someone manage to do this? They did it by flooding the servers, and their attached routers, with requests for information at a rate far too great for the system to handle.

Most routers have rules in the configuration table that won't allow millions of requests from the same sending address. If too many requests from one address are received in a short period of time, the router simply discards them without forwarding. The people responsible for the attacks knew this, so they illicitly planted programs on many different computers. These programs, when triggered, began sending thousands of requests a minute to one or more Web sites. The programs "spoofed" the IP address of the sender, placing a different false IP address on each packet so that the routers' security rules wouldn't be triggered.

When the **packet floods** were triggered, millions of requests for information began to hit the targeted Web sites. While the servers were being heavily taxed by the requests, the real impact was to the routers just "upstream" from the servers. Suddenly these routers, which were robust but of a size appropriate for normal traffic, were getting the levels of requests normally associated with Internet backbone routers. They couldn't handle the massive number of packets, and began discarding packets and sending status messages to other routers stating that the connection was full. As these messages cascaded through the routers leading to attacked servers, all paths to the servers were clogged, legitimate traffic

couldn't get through the logjam, and the attackers' goals were accomplished.

Web content providers and router companies have placed new rules designed to prevent such an attack in the configuration tables, and the companies and universities whose computers were used to launch the attacks have worked to prevent their systems being used maliciously. Whether their defenses, or the new attacks designed by criminals, will prevail remains to be seen.

Backbone of the Internet

In order to handle all the users of even a large private network, millions and millions of traffic packets must be sent at the same time. Some of the largest routers are made by Cisco Systems, Inc., a company that specializes in networking hardware. Cisco's **Gigabit Switch Router 12000** series of routers is the sort of equipment that is used on the backbone of the Internet. These routers use the same sort of design as some of the most powerful supercomputers in the world, a design that ties many different processors together with a series of extremely fast switches. The 12000 series uses 200-MHz MIPS R5000 processors, the same type of processor used in the workstations that generate much of the computer animation and special effects used in movies. The largest model in the 12000 series, the 12016, uses a series of switches that can handle up to **320 billion bits of information per second** and, when fully loaded with boards, move as many as **60 million packets of data every second**. Beyond the computing power of the processors, these routers can handle so much information because they are very highly specialized. Relieved of the burden of displaying 3-D graphics and waiting for mouse input, modern processors and software can cope with amazing amounts of information.

Even with the computing power available in a very large router, how does it know which of the many possibilities for outbound connection a particular packet should take? The answer lies back in the configuration table. The router will scan the destination address and match that IP address against rules in the configuration table. The rules will say that packets in a particular group of addresses (a group that may be large or small, depending on precisely where the router is) should go in a specific direction. Next the router will check the performance of the primary connection in that direction against another set of rules. If the performance of the connection is good enough, the packet is sent, and the next packet handled. If the connection is not performing up to expected parameters, then an alternate is chosen and checked. Finally, a connection will be found with the best performance at a given moment, and the packet will be sent on its way. All of this happens in a tiny fraction of a second, and this activity goes on millions of times a second, around the world, 24 hours every day.

Knowing where and how to send a message is the most important job of a router. Some simple routers do this and nothing more. Other routers add additional functions to the jobs they perform. Rules about where messages from inside a company may be sent and from which companies messages are accepted can be applied to some routers. Others may have rules that help minimize the damage from "denial of service" attacks. The one constant is that modern networks, including the Internet, could not exist without the router.