The view from the top left corner

Your Networked Computer: Checking Under the Hood

by Michael Schuyler

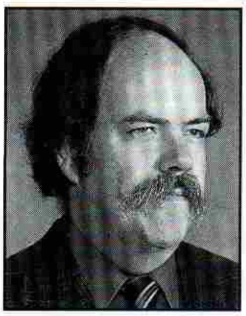

We use the Internet as easily as we use our cars. But what's really under the hood?

For those of you who detest "Information Highway" analogies, put this article down now. You can pick it up again in a couple of paragraphs.

I'm about to tell you why using the Internet is like using a car. You love this already, right? When you drive a car, you point it down the road and expect it to go there with very little attention on your part. I'll bet the only two gauges you ever look at are the speedometer and the gas gauge. There's an idiot light for oil pressure, and the only time you'll look at the temperature gauge is if steam is already pouring out from under the hood.

I'm teaching Linnea how to drive. I told her she could drive alone if she could take me for a spin and not scare me to death. Before we left on our last outing she told me I had better kiss Carol before we left because I might never see her again. When we finally made it home an hour later she had to open the door for me. I fell out gasping, kissing the ground and whimpering. She said, "Dad? Are you all right?"

After 35 years of driving I don't think much about it any more (though that 2,300 miles in England earlier this year was kind of inspiring), but this experience has renewed my interest in the fundamentals.

After only a few years we all use the Internet just as easily as we do our cars. When you press "Send" you expect your message to be delivered in a few seconds, end of story. So in this column we will return to the fundamentals of: TCP/IP! Yes, there will be some math, and some esoterica, but I promise that when you get to the end there will be a little surprise waiting for you.

The Joys of Using Protocols: TCP/IP

TCP/IP stands for Transmission Control Protocol/Internet Protocol, two of dozens of protocols that make up the "Internet Protocol Suite." TCP/IP is the part that ensures that data from one place gets to its intended destination intact.

To deliver a package anywhere in the world, the intended recipient needs to have an address. That's what the Internet Protocol (IP) provides, a unique address in the midst of the world's millions of computers, for every Internet-connected computer has an address. Every router has an address. Many printers, hubs, and CSU/DSU's have addresses. Even copy machines and fax machines may have Internet addresses.

Because of the way the Internet world was designed, each address must fit inside 4 bytes. That's 32 on/off switches, and not one extra. Simple math (2,564) gives you a possibility of over 4 billion addresses (4,294,967,296). Fine. There couldn't be that number of computers in the world, could there? That's one address per person, or will be shortly. (But actually some overhead considerations bring this figure down to something like 3.7 billion, still very likely enough for a while.)

The problem occurs in the way these networks are divided and assigned. There are three classes of networks: Class A,

Table 1: Network Classes					
	ID Bits	Net Bits	Host Bits	Nets	Hosts
Class A	1	7	24	127	16 million
Class B	2	14	16	16,384	65,384
Class C	3	21	8	2,097,151	254

Table 2: Traditional Sub-Nets		
	Typical Net #	Sub-Net
Class A	42.124.56.201	255.0.0.0
Class B	130.23.21.76	255.255.0.0
Class C	198.187.135.21	255.255.255.0

Class B, and Class C. There are only 32 bits available, divided differently depending on the class. The table above (Table 1) illustrates this.

In each class above, the identification, host, and network bits add up to 32. The smallest number of hosts that can be allocated with a network is 254. (Don't use 0 or 255 in calculations.) We have 10 Class C networks. Each one can handle 254 devices. At the largest library we have to be careful because we're using them all up, but at the smallest library we are using about a dozen addresses, leaving 240 or so lying fallow. Of the total number of addresses allocated to us—over 2,500—we're using a little over 300 of them, not much over 10 percent.

When we received our addresses, about 5 years ago, things were more lenient. No one thought the Internet would grow as fast as it has. If you have lots of water, you waste some. We had lots of Internet addresses, so we (meaning the country and the world) wasted some. We can't just "give them back" because this slop is built into the system and affects millions of Class C address holders. It also would mean a complete redesign of our network, including re-programming routers. The slop is far worse at the B and A levels. Indeed, there aren't any more of those to give out. I'm not sure we could get 10 Class C networks these days. They'd tell us to "sub-net" one or two of them.

Sub-netting is another wrinkle in this IP addressing scheme. We won't go into the specifics of sub-netting other than to say it is possible to sub-divide a network into two or more "sub-nets" by changing the number of bits associated with the host and with the network. Table 2 shows you the traditional sub-nets associated with each kind of network.

The number "255" is the largest number that can fit into a byte. There are 256 spaces, but 0 takes up one of those spaces. The sub-net simply expresses what kind of network this is. It is important because you can change it to create different sorts of networks. For example, you can take a Class C network with 254 host possibilities (remember 0 and 255 don't count) and "borrow" part of the hosts' numbering to add to the network numbering. Look at Table 3 to see what happens.

You can see what is happening a little better here if you look at sub-nets in terms of the binary numbering system. Table 4 shows the sub-net mask in base 10, then again in binary. You can see how a bit robbed from the hosts winds up in the network table, thus reducing the number of potential hosts.

Though we can "steal bits" from the hosts and add them to the network, there is a cost of smaller networks. If you steal 3 bits, for example, one network can have only about 30 addresses. For a really tiny branch that may be enough, but you must allow for growth. Using sub-net masking may help for a while, but it's not an ultimate solution to the dearth of Internet addresses.

Byting Off a Bit More Than We Can Chew

Well, you say, just add a byte and keep going! Sorry. It's not that simple. The IP addresses we have been discussing are just a part of a larger scheme called an "IP Packet" or a "TCP Packet," which includes "headers" that have space for various information. They have a space for the source address. It's 32 bits long. They have a space for the destination address. That's 32 bits long. To add a byte would not only destroy the symmetry of the header; it wouldn't work unless the programs that make the Internet work were completely

Table 3: Sub-Net Masking			
	Sub-Net Mask	Borrow a Bit	Borrow 2 Bits
Class C Network	255.255.255.0	255.255.255.128	255.255.255.192
Number of Nets	1	2	4
Number of Hosts	254	126	62

Table 4: Sub-Nets Compared to Binary: "Stealing Bits"	
Sub-Net Mask for Class C Network	1's Are Network, 0's Are Hosts
255.255.255.0	11111111.11111111.11111111.00000000
255.255.255.128	11111111.11111111.11111111.10000000
255.255.255.192	11111111.11111111.11111111.11000000

> ***"Memory was expensive back then, so it paid to save a byte here and there."***

re-designed. That means every single encapsulated program in every router, every CSU/DSU, every Network Interface Card (one per computer), every program, every piece of equipment hooked to the Internet, would have to be changed. Many of the programs are embedded in the equipment itself in "PROMs" (Programmable Read Only Memorys), so you can't just boot up a floppy and make the change.

Still, we're further along. We know that a complete Internet address is 4 bytes long. We know that a byte can hold numbers up to 256 (28). We also know that the network identification bits on a Class C network must be set to 110. Therefore, we know that every Class C network on the planet must start with the number 192 or higher because otherwise those bits and bytes won't work out right.

We're left with a 4-byte notation, called "dotted decimal" to make up a network. You can tell by looking at the first number what kind of address this is. Table 5 shows this in practice.

Every computer hooked to the Internet has an IP address, either permanently or temporarily assigned. Of course, it doesn't mean much to have a string of numbers assigned to each computer, so something called the "Domain Naming System" was invented to lend a little character to all this. This provides for more-or-less discernable names to be associated with the IP numbers. The translation is provided by a Domain Name Server (DNS). You either run a DNS server yourself or you pay someone to do it for you. Table 6 shows a few of the translations for our networks.

So when you type something like: http://www.kitsap.lib.wa.us, your DNS translates that to http://198.187.135.222 because your router, by virtue of being connected to the Internet, has had this address translation "propagated" to it over time. If we make a change to the machine and start using another one, we'll tell our router, which will tell its neighbor, which will tell its neighbor until eventually your neighbor will tell your router that we changed the number, from a few hours to a few days ago, depending on how efficiently everything gets routed.

The router is a specialized computer that helps hook you to the Internet. It advertises its presence and acts as a traffic cop, sending your packets out into the Internet and allowing only those addressed to your network back in. The router's own IP address is called a "gateway," which makes a certain amount of sense. Your computer is told about the gateway when it is set up, so it knows to send packets to the gateway to be routed out. There is no hard rule about what number to use: We've settled on #1 of any of our networks as the default gateway, thus the gateway for our central network is 198.187.135.1. You could just as well call this 123; it's just that you need to keep it all straight. All but one of our routers are getting kind of old. They're based on 386SX computers with about 2 megabytes of memory—dinosaurs, really, but they continue to do the job just fine.

> ***"Though we can 'steal bits' from the hosts and add them to the network, there is a cost of smaller networks."***

Addressing, Coding, and Translating Packets

That's all fine, but how does the Internet know your computer is yours? You have an IP address assigned to your computer, but that's arbitrary. You or I could make the decision as to which IP address to use. How does the Internet know a given packet is intended for your physical machine?

The answer is in the MAC address. MAC stands for Media Access Control and is a string of numbers hard-coded into the Network Interface Card plugged into your computer. A typical MAC looks like this: FA-08-12-34-C0-A1. This is "hexadecimal coding," which just happens to be an easy way to represent an 8-bit byte. The first 3 bytes in this code are owned by a company. The second 3 bytes are a serial number. In this case there are spaces for 16 million companies each with 16 million serial numbers, so at least this part of the overall Internet scheme won't have to be redesigned any time soon.

Your router picks up these MAC addresses and relates them to the IP address assigned by your network administrator. It is the router that translates between the two and knows where to send packets addressed to your machine.

Table 5: IP Address Construction

	Network Part	Host Part	Whole
Class A	27	123.64.231	27.123.64.231
Class B	130.119	76.213	130.119.76.213
Class C	198.187.135	22	198.187.135.22

Table 6: IP Numbers and Domain Names

Internet Address	Domain Names
198.187.135.22	Deepthought.kitsap.lib.wa.us
198.187.135.21	Linknet.kitsap.lib.wa.us
198.187.135.222	www.kitsap.lib.wa.us
198.187.135.16	Webpac.kitsap.lib.wa.us

Table 7: Internet Requirements

Requirement	Example
A unique IP address	198.187.135.22
A defined sub-net mask	255.255.255.0
A unique Network Interface Card	00-60-08-99-65-5C
A defined gateway	198.187.135.1
A name	Linknet.kitsap.lib.wa.us
A designated DNS server	198.187.135.18

So finally we have all the components we need to hook your computer to the Internet. (See Table 7.)

Perhaps you've never seen any of this, but it's there under the hood of your computer somewhere, nevertheless. It doesn't matter whether you're using character-based PCs or the latest in Windows NT technology, somewhere there is a file inside your box that tells all this information to the Internet.

Windows 95/98 is particularly easy in this regard. You can create an entire network just using Windows 95. In fact, I am beginning to wonder why we even need expensive file servers any more, since networking is so easy with this product. We recently upgraded our Training Center computers by removing the Novell file server entirely and turning it into a strictly Win95 network.

That begs the question of this Internet address problem, but I'm not sure we'll have to worry about it. After all, good minds are working on a solution right now. What minds? Good minds. There are also good minds working on another minor problem facing us.

You know now that you can fit any number up to 255 in a byte. You know a byte is a fundamental building block of our computer infrastructure, deeply embedded in everything, often hard-coded inside computer chips, particularly those used in such things as routers and process controls used in industry.

255 is a lot of space. Well, it was! You can fit every letter of Western European languages in 255 spaces, numbers, a lot of diacritics, more than enough space for that sort of thing. So much, in fact, that IBM came up with its own character set and stuck corners of boxes and line drawing characters in the "upper 128" spaces so you could draw tables and graphs in character-based DOS.

Now You're Ready for Your Little Surprise

You've endured a lot, so here's the surprise: You can fit the number "98" in a byte. But you can't fit the number "1998" in a byte. So all the programs written in ancient programming languages such as Cobol, those programs that control a lot of the financial industry, probably have a lot of 1-byte dates in them. Memory was expensive back then, so it paid to save a byte here and there. That's this "y2k" problem you've been reading about.

Some of these 1-byte dates are built into process control chips the same way TCP/IP routing protocols are built into routers. Some of those process control chips control such things as electricity generating plants. They control such things as maintenance schedules. So, when the clock ticks over into what really is the last year of our millennium, not the first year of the next, all those chips will calculate 00 from 99 and think maintenance has not been performed in over 99 years. The question is: Will they all shut down?

Have a nice day.☺ ▲

Michael Schuyler is the systems librarian for the Kitsap Regional Library System in Bremerton, Washington. His e-mail address is michael@linknet.kitsap.lib.wa.us.

Look what people are saying about
Marketing Library Services

The "How-To" Marketing Tool Written Specifically for Librarians!
"If our clients read MLS as diligently as we do, there would be little need to hire us as consultants. Kudos to your subject mix and timely presentations. Shh…don't tell our customers."
Gloria Dinerman, President
The Library Co-Op, Inc.
Edison, NJ
"I never read an issue without taking something away from it."
Carol L. Ginsburg, Vice President
Bankers Trust Company
New York City
TRY IT YOURSELF AND SEE
For a FREE TRIAL ISSUE of *MLS* simply call **1-800-300-9868** or fill out the form below and send it to:
Information Today, Inc.
143 Old Marlton Pike
Medford, NJ 08055
☑ YES! Please send the latest issue of *Marketing Library Services* absolutely FREE. If I like it, I'll pay just $69.95 for 7 additional issues (8 in all).
NO COST. NO OBLIGATION.
If I'm not completely satisfied with my preview copy, I'll simply write "cancel" on my invoice and return it. I'll owe nothing at all, and the free issue will remain mine to keep.
Name ____________
Title ____________
Organization ____________
Address ____________
City ____________
State & ZIP ____________
Phone Number ____________
ITI

Copyright of Computers in Libraries is the property of Information Today Inc. and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.