SWCP PORTAL
Providing News, Views, & Clues from Your Local Internet Experts **201007**

As if we didn't have enough already...

Addressing the Net Address "Crisis"

By Jay Nelson, editor

The Internet is big, *really* big, and growing at breakneck speed, too. With millions of machines chatting back and forth across the planet, the globe is awash in oceans of electronic traffic. How is it that everything from tweets to video streams gets to its proper destination?

Like everything else that must be located, **addresses** are the key. Each and every gizmo on the Internet, all those PCs, laptops, smartphones, tablets, routers, and servers, *must* have its own unique and locatable address. They must be matched up with all the websites and locations of other services in cyberspace, too.

The clever solutions to this intricate and daunting problem shaped the Internet. Unfortunately, due to the unforeseen, explosive growth of the Net, a key part of the current system, called "**IPv4**", is already utterly inadequate. This means that the Internet is running out of addresses – perhaps even by the end of this year.

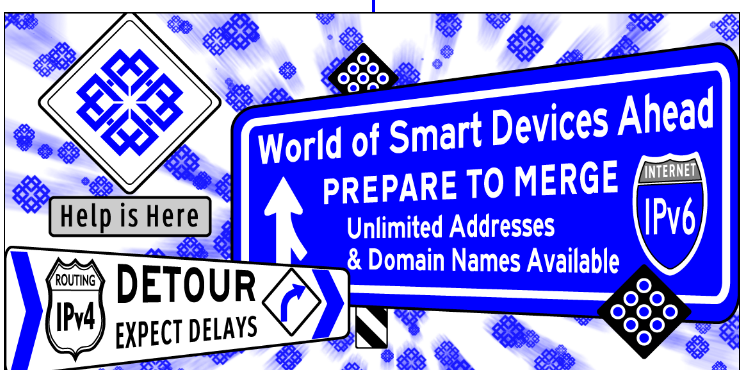

A new method, "**IPv6**", must be employed soon. Just how quickly, no one knows, and the changeover may not be simple or trouble-free. The standards have been agreed upon and the transition is already underway, but the Net might become crowded, slow, with erratically unreachable sites for a long time before it's over. Or not. The problem with predicting just how it will go is that nothing like this has ever happened before, so nobody really knows how difficult the conversion will be.

Once accomplished, however, the Internet will open up and transform our planet like nothing has before. Imagine a smart world where *everything* can be connected – not just information appliances but everyday objects like refrigerators and medicine cabinets (and all their contents), washing machines, table lamps, air conditioners, toasters, clothes, furniture, etc., etc. IPv6 will enable a full-blown, world-encompassing "**Internet of Things**" linking intelligent objects whose potential we can barely imagine. This article will touch on how and why this change is going to happen and the effect on everyday users this momentous changeover will have.

Telling machines apart

Once computers are connected, they must be distinguished from each other. In 1969, the developers of **ARPANET**, the Defense Department's experimental network, wanted to link mainframes at different universities to share computing power on big scientific research problems. They wanted an open-ended system that could both allow the addition of more machines easily and indefinitely yet be rigorous enough to make identifying each one simple.

Names were okay for direct connections. But the invention of **packet switching** to allow many computers to communicate at the same time made a new system necessary. Every small package of data passed from one computer to another has to have an address that all machines could use.

IP (Internet Protocol) addresses were the answer, a unique one assigned to each machine. Composed of four blocks of three digits each, the first parts basically represented the network of computers and the last specified the individual device or gateway.

It was an elegant solution but quickly used up addresses. Techniques were developed to extend and reuse IP addresses for private networks. The main one was **NAT** (Network Address Translation) which hides private local networks behind a single IP address gateway to a larger, public network.

NAT quickly became indispensable for routers for homes and small businesses. It conserves addresses, but NAT fudges the originally envisaged wide-open end-to-end connectivity of the Internet, and raises hosting and performance issues. Another technique, **DHCP** (Dynamic Host Configuration Protocol) assigns IP numbers from a pool, and so is best suited for intermittent connections like dial-up rather than always-on broadband.

What's in a name?

But numbers solve only half the addressing problem. They merely identify devices, not resources. Whereas

Continued on back

SOUTHWEST CYBERPORT **5021 Indian School NE, Suite 600, Albuquerque, NM 87110-8910 USA**
help@swcp.com | Phone: (505) 232-7992 | swcp.com | *Member, AIBA* keepitquerque.org

Continued from front

machines depend entirely on numbers, people find names much easier to recall. Therefore to match up content to computers, the **DNS** (Domain Naming System) was created. These are the addresses every Web user is familiar with, such as SWCP's own **URL** (Uniform Resource Locator): *www.swcp.com*.

Like IP numbers, each section represents a smaller portion of the Internet, but read backwards instead of forwards. The last part, "com", for commercial, is a **Top Level Domain**, denoting the kind of institution we are. Originally, they were very limited in availability: along with "**com**", there were "**net**" for network, "**edu**" for educational institutions, "**org**" for non-profit organizations, "**gov**" for government agencies, "**mil**" for the US military. **Country Codes** could be used for places outside the USA, like "ca" for Canada, and "tv" which is actually for the tiny island nation of Tuvalu, not television.

The middle part, "swcp", is our own unique **Domain Name**. No one else can use it for a commercial site. The first part, "www" for "world wide web" is for our Web pages, as "ftp" would be for our file server. Called a **subdomain**, some websites don't even use them.

Of course, someone must keep track of all this data. Every one of these domain names and IP numbers in use anywhere is governed by **IANA** (Internet Assigned Numbers Authority). IANA has recently taken steps to vastly enlarge the pool of Top Level domains to cover sites from "eu" for the European Union to "mobi" for mobile devices, and special designer brands like "disney". Even the sex industry will be served, as "xxx" for porn sites has been approved as well. Whether this will in any way serve to limit and isolate the spread of Internet pornography as it is intended remains to be seen.

The Net authority has also allowed the use of non-Latin alphabets. In so doing, IANA has ensured that names at least will never be exhausted. The lists of active domains, like a huge phone book shared by 13 giant "**root name servers**" across the globe, are constantly updated so they can be accessed by servers everywhere to look up addresses.

The good news is that not only are domain names now practically unlimited but the DNS system will be unaffected by the IPv6 changeover. All your bookmarks and links should still work just fine.

Crisis by the numbers

There are "only" 4.3 billion IPv4 numbers available – far less than the current world human population. Since IPv6 numbers are *much* longer, there is an astronomically insanely larger amount of IPv6 addresses possible— too great to be graphically compared. But there are *almost twice as many IPv6 addresses as stars that exist in the whole Universe* – quite enough to literally label *everything* on the planet with its own unique IP number if desired.

To give another idea of just how mind-bogglingly *many* IPv6 numbers there are, SWCP has already been assigned the smallest block of the new numbers available to ISPs — and we've been granted more than the world's *entire* IPv4 space!

However, a real crunch is coming shortly. Due to NAT and other tricks, neither day nor hour can be determined, but IPv4 address exhaustion will occur all around the globe. There is no fixed turnover date like Y2K to serve as a deadline, and filling up unused addresses may ease the predicament some, but there will surely come a point when all IPv4 addresses are taken.

How does this become a crisis? Simple: **IPv4 and IPv6 are *not* compatible**! Besides the longer sequences, changes in size, formatting, and handling of IPv6 packets means that the software for IPv4 will not work with them without fixes, and vice versa. While most recent operating systems can run both systems in parallel, software patches for some web browsing, email, IM, FTP, and other applications might not be ready in time.

Implementation of the IPv6 protocols is seriously lagging everywhere. Few institutions and even fewer users are in any way prepared or even aware of the change. There will likely be little demand for IPv6 until all IPv4 addresses are gone — but then, watch out.

At the moment, since IPv6-enabled sites are still quite rare, there are few other IPv6-ready devices to contact. Not only that, but there's nothing but IPv4 devices between them. To connect, new techniques, such as repackaging IPv6 packets in IPv4 wrappers, must be worked out for various platforms, debugged, and installed across the entire Internet.

So the transition looks like a long, unpredictable, and perhaps bumpy ride to our glorious promised future.

Southwest Cyberport is already working on our turnover plans. We will stay on top of the situation, doing our best to make the changeover smooth and painless for all of our customers. In future issues, we'll examine just what IPv6 will do for the Net. Meanwhile, count on your Internet experts, SWCP, to remain on the cutting edge and to keep you abreast. Stay tuned!

Net Notes

Number games

Like everything else on the Net, both IPv4 and IPv6 addresses use binary digits, but they are expressed differently — IPv4 in octets or bytes of 8 bits in 4 groups separated by periods, and IPv6 in hexadecimal format (base 16) which uses 0-9 and A, B, C, D, E, and F as well, in groups of 4 separated by colons. Got it?

IPv4 addresses are **32 bits** or 4 bytes in length, and appear in a format like this: **192.168.8.1**

IPv6 addresses are **128 bits** or 16 bytes long. The much greater length makes them easily recognizable from IPv4 addresses, and may look something like this: **2001:0db8:85a3:08d3:1319:8a2e:0370:7334**

Total number of IPv4 addresses is: **4,294,967,296**

And the total amount of possible IPv6 addresses is: **340,282,366,920,938,463,463,374,607,431,768,211,456** — which should take care of the problem!