STILL WAITING FOR INTERNET PROTOCOL VERSION 6

The Internet Protocol version 4 began in 1981, and eventually someone will send the last packet using it. I'll bet you they will send it from a horse-and-buggy North American ISP.

Sixteen years ago, the replacement for the Internet Protocol was made official. Years of work by groups of experts from around the world had solved the problems that were faced by the old way of doing things, which was called Internet Protocol version 4. For a short while, they called this grand collaborative effort IP "Next Generation" — lots of Star Trek fans in that bunch! — but eventually it was decided to call the new, improved set of rules "Internet Protocol version 6," or IPv6 for short.

What is wrong with IP version 4? Nothing at all, if you think about the Internet as an experimental and educational network of networks — which is exactly what it was in the beginning. When Vint Cerf and Bob Kahn and others working on networking principles thought about how many addresses would be needed, deciding to use a number that was 32 bits in length seemed like more than enough. Having more than four and a quarter billion possibilities when there were only a couple of hundred nodes in the ARPAnet (the early Internet) was dreaming beyond possibility, to these pioneers. Other issues like security and efficiency of routing and autoconfiguration never really came up. Everyone using TCP/IP networking was an expert, and IPv4 could be adapted to meet requirements until something better came along.

What came along was the World Wide Web, and the Internet as a business, and eventually as a communications medium for everybody. Suddenly the set of rules for identifying computers on the network, and for moving packets of data from one of them to another, was facing severe strain. The most obvious problem was that the set of available numbers was going to run out. Briefly, an Internet Protocol address needs to identify both the computer and also the network it is in. Portions of the overall address space are handed out to people who administer networks, and from that "piece of the pie" administrators can identify individual computers that are on the network they are responsible for.

A pie can only be divided up so many times, though. In order to keep IPv4 going, several life extension strategies were derived for it. The three that were most successful were subnetting, dynamic host configuration, and network address translation. Subnetting can be thought of as dividing up the pie as efficiently as possible. Instead of giving administrators fixed network sizes, why not give them the size that is closest to their needs, and then allow them to sub-divide that on their own if they wish? Subnetting dealt with the wastage of address space that was happening with the older, inefficient way of dividing up the pie. Dynamic host configuration is the simple idea of giving a computer an address only for the time that it is actually connected to the network, and having it give up that address up when it is not. With this strategy, a small pool of addresses could serve a large number of computers, as long as they are not all connected at the same time. Network address translation involves giving a number of computers private or arbitrary addresses, and having many such addresses switched to one "real" address in order to communicate on the Internet. A network address translator device is usually a firewall, or a router connected to the Internet. Such a device needs to keep track of the mapping between the fake internal addresses and the real external address, and switch back and forth with every packet of data that flows through it. With such a device, a large number of computers can communicate on the Internet as if they were individual nodes, when in fact they will appear to the outside world to be one (very busy!) machine.

All of these life extension strategies have reached a breaking point. Two years ago, the Internet Assigned Numbers Authority gave out the last piece of the pie. Technically, they assigned the last /8 CIDR block to one of the large geographical registries. There will never be a new IPv4 address again, and we must re-use the ones that we have got to get by.There is no need for worry, though, because IPv6 is available right now, fully-tested and functional, and it makes IPv4 addressing problems disappear. IPv6 has a 128 bit address space. That's 340 undecillion possibilities. This allows for a staggeringly large number of networks to exist, every one of which will have a staggeringly large number of uniquely identifiable real host identities on the Internet. Effectively, an Internet Protocol address has become as free as air. Strictly speaking, air in the atmosphere is finite, but it is large enough in supply to freely support everyone who will ever live on this planet.

IPv6 makes network address translators obsolete. That bottleneck is gone. In fact, a system of simplified headers, and building functionality like autoconfiguration into the Network Layer instead of the Application Layer, means that moving packets of data is now much more efficient overall. With the artificial scarcity of addresses gone, ISPs can focus on the business of providing a fast and reliable connection to the network — which is the business they should be in.

It would seem that common sense and self-interest would be driving ISPs to rush to IPv6, to run their businesses as efficiently as possible. This has not been happening. As a whole, North American ISPs have paid scant attention to IPv6, and have put their efforts into the life extension schemes for IPv4. It has gotten so bad, that Vint Cerf has warned about the emerging phenomenon of “cascading NATs.” This is where one network address translator is not enough to hide the large numbers of computers, tablets and smartphones needing to connect to the Internet, and so a series of network translators is configured behind another network address translator. A large number of fake addresses are hidden behind a smaller range of fake addresses, which in turn are hidden by the one, precious IP address that is available. Imagine the comparison of plugging a power strip into an outlet in the wall, and then plugging a series of power strips into each outlet available on that power strip. Will it work, and will you be able to power up a large number of electrical devices? Up to a point … and then the circuit breaker will trip and all the devices will fail. This is what is happening with cascading NATs and IPv4. It is a bottleneck on top of a bottleneck, and it is madness. It is also what many North American ISPs are doing, instead of switching to IPv6.

North American ISPs are living in a fool’s paradise, because of the disproportionately large portion of IPv4 address space that was originally allocated to this geographical registry. But it will not be a paradise for much longer. Places like Japan and Europe faced the address space crunch earlier, and ISPs in those places give out IPv6 networks to customers and run IPv6 routers, and are already living in the new reality that we will all inhabit eventually.North American ISPs are destined to be the Western Union of the Internet. In 2006, someone sent the last telegram from a Western Union office, and that was the end of a mode of communication that began in 1844. The Internet Protocol version 4 began in 1981, and eventually someone will send the last packet using it. I’ll bet you they will send it from a horse-and-buggy North American ISP.

Source : https://www.exitcertified.com/blog/michael/2012/still-waiting-ipv6/