Nigel Titley

IPV6 AND THE INTERNET ADDRESS CRISIS

You may or may not have heard of Internet Protocol version 6 (IPv6), says Nigel Titley. Here he clarifies what it is and how it differs from Internet Protocol version 4 (IPv4), as well as explaining why he believes that it is imperative to migrate the Internet from v4 to v6

IPv4 was specified in RFC791 [1] (an RFC is a Request For Comments, the fundamental standards specification document in the Internet world) which was released over 30 years ago in September 1981 and has stood the test of time, still being the main protocol in use on the Internet. The Internet protocol is packet-based, that is to say data is transported across it as a series of discrete packets. This article is only concerned with the packet layer and not how packets are assembled into data streams. Figure 1 shows the layout of the IPv4 packet header.

Most of these header fields are self-evident but a few are worth looking at, in particular the Fragment Offset and the Source IP Address and Destination IP Address fields. It is also worth noting that there may be an indeterminate number of options, so the header length cannot be predicted. In general and at a very high level, an IP packet is injected into the Internet by an end-device (such as a PC) and then passed through the network from one intermediate device (usually referred to as a router) to another until finally arriving at its destination. At each intermediate device a decision is made on how to forward on the packet.

IP packets can be carried over many different media; in fact this is one of the strengths of the Internet, and at each change of media (for example from Ethernet to Synchronous Digital Hierarchy) the router will perform the necessary repackaging of the packet. Different media may have different maximum packet sizes and a router shifting an IPv4 packet may need to split or "fragment" it when moving it from a medium that allows a long packet length to one that only allows a shorter one. This is the function of the Fragment Offset field which shows that the packet being received is a fragment of an originally larger packet and is a candidate for reassembly.

The address fields are what this

article will mainly discuss. As you can see from Figure 1 these are four octets, or 32 bits long, allowing a total of approximately 4 billion unique addresses. They are generally written as what is called a dotted quad, for example 192.168.1.52 where each quad is an 8-bit byte expressed as a decimal number.

In 1981 the idea of a network with 4 billion end devices was unthinkable. In fact originally these address fields were set at 8 bits long, allowing for a paltry 255 unique end systems, which in the days of mainframes and with a network planned only to encompass North America, seemed ample! Thankfully, wild optimism prevailed and 32 bits was chosen. However, for those who have somehow missed the excitement and warnings over the past five years, the supply of these addresses is about to run out (indeed has run out in certain regions). The replacement protocol IPv6 was designed and formalised starting with RFC1883 [2] in 1996 (17 years ago) but has not, as yet, got a great deal of traction for a variety of reasons.

IPv6 has a number of improvements over IPv4. Many of these are intended to reduce the processing that routers have to do as they move packets around. Figure 2 shows an overview of the IPv6 header layout – and it looks much simpler. Note that the Fragment Offset field has gone as IPv6 routers never fragment packets. The IPv6 protocol uses a technique called MTU (Maximum Transmission Unit) path discovery

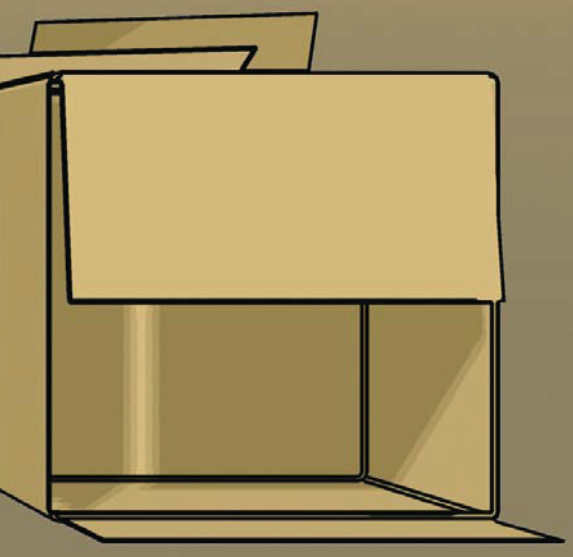

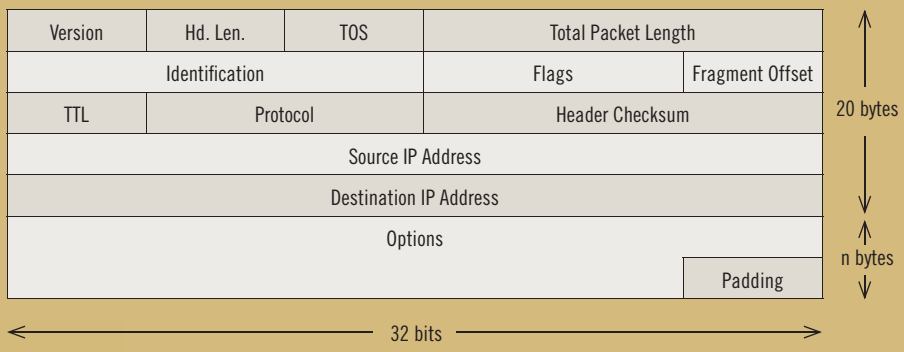

Figure 1: *Layout of the IPv4 packet header*

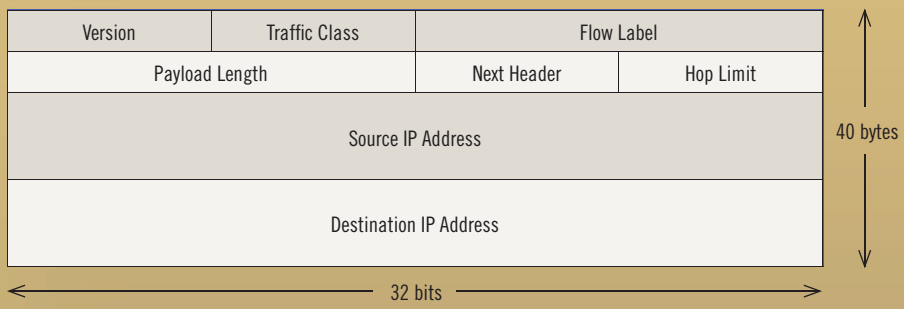

Figure 2: *Layout of the IPv6 packet header*

to find out the maximum size packet that the path to the other end of its connection can carry and an end system should never send packets that are larger than this. Also, the header is now a fixed length, making for hugely improved ease of processing at each intermediate device. However the important thing to notice is the size of the address fields. They are now a huge 128 bits long, enough for roughly 400 trillion-trillion-trillion unique addresses. IPv6 addresses are written as 32 hexadecimal nibbles (a nibble is half a byte). An example is: 2a00:1940:0000:0000:0213:8fff:feb9:057d

These are obviously more difficult to remember than IPv4 addresses and some simplification is allowed: leading zeros in each 16-bit group (or "chazwazza") may be removed and multiple all-zero chazwazzas may be collapsed into a single double colon. Thus the above address may be simplified to: 2a00:1940::0213:8fff:feb9:057d which helps somewhat.

IPv6 should solve the problem of lack of Internet addresses for a long time if only it can gain some traction. Unfortunately most of the larger Internet Service Providers (ISPs) have been very slow in rolling it out. This is partly due to the lack of available low-cost consumer access devices and partly due to economic forces: there have been no long term business case drivers to install IPv6.

In practice most Internet connections consist of a data provider and a data consumer. There are many, many exceptions to this, of course, but at the moment the typical model is a commercial provider of data, such as a web site or video streaming provider and a consumer, typically using a web browser or smartphone application.

The vast majority of consumers are home broadband customers, using a modem supplied by their ISP. A large ISP will be providing many millions of these modems and a difference of a few pence in their price will save substantial amounts of money. The manufacturers are hence competing on price and will cut the capability down to the bone. One obvious feature not to include is IPv6, with its requirements for a completely different protocol stack, increased memory and higher powered processor. For those customers who buy their own modem, price is likewise a large driver. The net result is that, although many ISPs

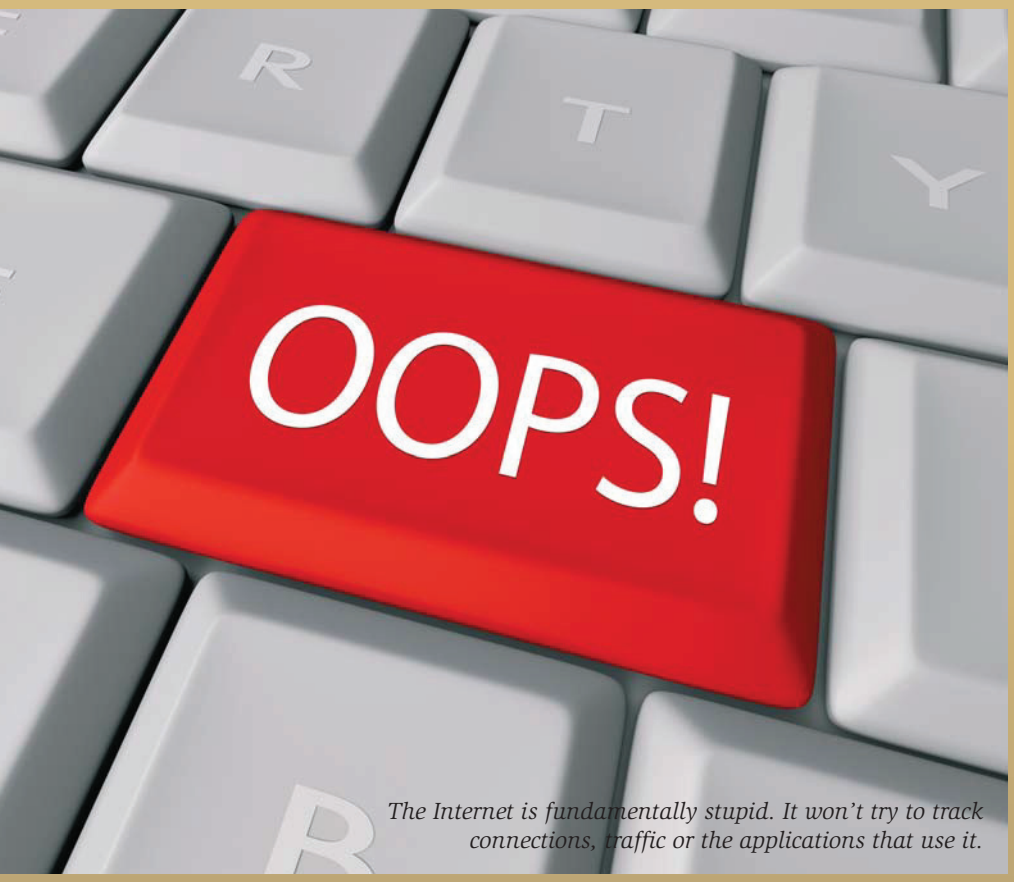

The Internet is fundamentally stupid. It won't try to track connections, traffic or the applications that use it.

have infrastructure that is capable of delivering IPv6, their customers have nothing to terminate it on.

There are signs that this is now gradually changing and middle-end modems are starting to include IPv6 at last.

Now for the other end of the connection. The data supplier generally has little difficulty in adding dual-stack (IPv4 and IPv6 running in parallel) to their product but sees no demand from the consumer, who in general doesn't worry about the underlying protocol stack as long as they can see the latest blockbuster film. If it weren't for some visionary data suppliers, such as Google and Facebook, IPv6 traffic levels would be much lower than they in fact are.

There are other factors that have slowed the uptake of IPv6: fear by Information Technology (IT) professionals, extra charges by ISPs and router vendors are typical examples. IT has become such an integral part of business today that IT workers become understandably nervous of new technology. IPv6 with its new address formats (far too long to be remembered and manipulated easily) and the fact that IPv4 seems to be working well (which it will, until we run out of addresses) tends to be swept under the carpet and ignored in favour of applying the latest Windows service patch which may be seen as more urgent. Firewalls are another sticking point. Many current internal networks use private addresses which are translated to real external IPv4 addresses by Network Address Translation (NAT) devices. NAT has been part of firewall design for so many years that security administrators are understandably loath to take it out, but once that has been done, the freedom to concentrate on the firewall rules, rather than the address mapping can be remarkably liberating and usually leads to far cleaner and easier to understand designs.

> "Any modern operating system is able to use both IPv4 and IPv6 simultaneously and users won't generally know the difference. The main changes are in the applications."

Finally, of course, manufacturers have seized the opportunity to charge extra for IPv6, even router vendors whose continued existence depends on a healthy Internet.

There are several possible workarounds for the shortage of IPv4 addresses and the most commonly suggested is the use of Carrier Grade NAT (CGN), which is a specialised router embedded within the network and which converts IPv4 to IPv6 and vice versa. CGN can help to spread out the available address space and make it last longer by allowing an IPv4 address to be used for multiple customers. CGN is a quick fix and may seem to provide a solution, but for one major problem, for which we need to go back to the reasons for the incredible success of the Internet.

The most important thing about the Internet, as opposed to some of the networks that preceded it, is that it is fundamentally stupid: it does not attempt to track connections, traffic or the applications that use it. This means that when a new application is designed by a teenager in their bedroom, the Internet can instantly transport their traffic. There is no need to negotiate with the Internet transit providers to allow this new application to be transported over their networks. All of the intelligence is in the network edge. It is this network stupidity that is behind the success of the Internet and it is why CGN is such a bad idea because a CGN has to be clever and make assumptions about the traffic flowing through it and these assumptions may be wrong.

Good ISPs recognise this and have been building their networks to transport IPv6 for many years. They will offer both IPv4 and IPv6 and in most cases their customers, once the idea has been explained to them, will happily use both protocols. Any modern operating system is able to use both IPv4 and IPv6 simultaneously and users won't generally know the difference. The main changes needed are in the applications, and even these are

generally limited to fixing user interfaces that assume that IP addresses are always 32 bits.

There are other side-effects of the shortage of addresses, ones that any first year student of economics could predict. Firstly, IPv4 addresses will start to be seen as assets. It is worth taking a closer look at this. IPv4 addresses come in two types: the first type is the kind that every commercial Internet user is probably used to dealing with. It is loaned to users by their ISP and returned when the user changes provider. For ISPs, these blocks of addresses are a vital resource. They cannot do business without them. These addresses originate with the central Internet custodian of number resources, IANA (the Internet Assigned Number Authority), are delegated downward to the various Regional Internet Registries (RIRs), are allocated to ISPs and are loaned or assigned by the ISPs to their customers.

Some of the RIRs have policies allowing the transfer of addresses to another user (ISP), and a transfer fee will usually be negotiated between the seller and the buyer. For large blocks of addresses these fees can be eye watering: a typical price is $10 per address. If a company has one of these blocks available and free then its Chief Financial Officer (CFO) might suddenly sit up and take interest (as might the tax man, as these have become an asset and subject to capital gains tax). If a company needs one of these blocks then its CFO will certainly take interest, because what used to be a zero-cost resource is suddenly going to cost a lot more.

> "What else can ISPs try? They can't buy addresses for long. It isn't possible to manufacture more of them...Even if all the addresses not visible on the Internet suddenly became available, the World would still run out of IPv4 in the next five years or so."

This price is not going to go down; addresses are a scarce resource and no more of them can be created. Indeed there are already companies being set up to help match up those with address space spare with those who desperately need it. And they take a fee for acting as the marriage broker.

So what will happen next? An ISP needs to connect a new customer and it doesn't have an IPv4 address for them. So it will roll out a CGN and start putting customers onto private address space (private address space is a number of address blocks defined by RFC1918 [3], which are probably already being using within a company or at home). Of course the ISP can only do this until private address space runs out for them and after that each block of private addresses won't be able to exchange traffic with the others, which has implications for peer-to-peer protocols like Skype. This approach looks like it might be able to buy a little time but it's a dangerous path to tread.

As ISPs roll it out they are making their little bit of the Internet less and less stupid, more and more opaque. The Internet starts trying to guess what the end users are trying to do and, although it is likely to guess right when they are just browsing a web site, it is likely to get it wrong if they try and do anything more complicated. It is particularly likely to get it wrong when a new application comes along that it hasn't heard of. The end users will eventually notice that their packets are being intercepted and readdressed and, just like those "opened by customs" stickers on parcels from abroad, they won't be sure that the contents haven't been tampered with. And of course, every time the packet is held back for inspection it will be delayed.

Eventually, of course, an application will come along that the CGN doesn't understand, and then the end users won't be able to use it. The ISP has already started to lose customers because they cannot give them a unique IP address and now they are starting to see poor latency and on top of that the latest Internet craze doesn't work.

So what can they do? How about a walled garden? Restrict the end-users to only the things they are allowed to see. Much as governments might like to see this, in effect the ISP is building a firewall around their network and also possibly putting application gateways in as well. This technique used to work for mobile phones and some early ISPs, but Internet users have got used to going everywhere and suddenly they find barriers and locked gates.

The global economy is completely dependent on the Internet

How long will such ISPs' market share survive?

What else can ISPs try? They can't buy addresses for long. It isn't possible to manufacture more of them. The rate of consumption of addresses is such that even if all the addresses not visible on the Internet were to suddenly become available (and don't forget that just because they aren't visible in the global routing tables, they may well still be in use), the world would still run out of IPv4 in the next five years or so.

Perhaps we can stop using the Internet altogether? No, that probably isn't an option at this stage in the game, where the global economy is completely dependent on the Internet and the near instant access it gives to, literally, a world of information. Our stock markets and banking systems depend on it. Our children use wikipedia to write their homework. We order our groceries using a supermarket web site or a smartphone app. Our businesses depend on it to deliver orders and invoices, quotes and contracts. Commerce is lubricated by the grease of the Internet and without it the wheels will start to squeak and the axles will run red-hot, and what was a formula-one racing car will become a horse and cart. Recall teleprinters and telex machines and the cost of a transatlantic calls at £5 a minute? All of that is a distant memory thanks to the Internet. No, we won't be getting rid of it just yet, at least not until there is something bigger and better.

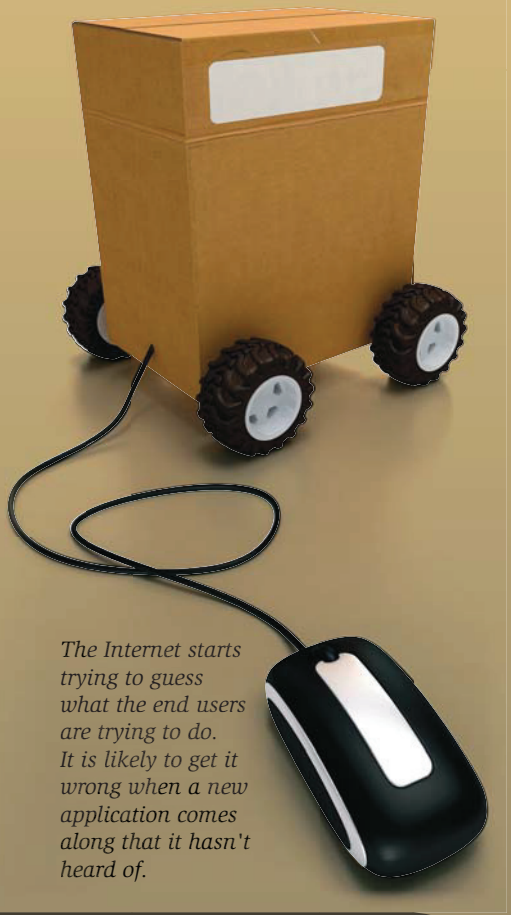

The Internet starts trying to guess what the end users are trying to do. It is likely to get it wrong when a new application comes along that it hasn't heard of.

AUTHOR'S CONCLUSION

It does begin to look like the only solution is IPv6. Most PCs and servers already work with it, once attached to an IPv6 enabled network. IT staff will complain, but with management support they will move forward. They will learn about it in the same way that they learn about the latest bug in Windows 7. Their firewall configurations will be simpler and more bug-free without the need to do the gymnastics necessary to get things like Voice over IP working through NAT.

Some ISPs will claim that they can't support it. But users can move to ISPs who can (there are plenty of them out there). Most Internet users at home and at work won't even notice. The IPv4 address marriage brokers reckon they have about four years of business model after which IPv4 addresses will have no value and the world will have moved to IPv6. The IPv4-only web site will have gone the way of the dinosaurs, the teleprinter and the ticker-tape machine.

Like it or not, IPv6 is coming at last. The exhaustion of IPv4 address space has been put off as long as it could be. Asia ran out of addresses in 2011. Europe ran out at the end of 2012. North and South America will probably hang on for another two years and Africa has got until 2018. Africa oddly enough, with less legacy routing equipment is probably best placed to surge forward into IPv6, making the exhaustion of IPv4 a truly first-world problem.

ABBREVIATIONS

CGN	Carrier Grade NAT
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISP	Internet Service Provider
NAT	Network Address Translation
RIR	Regional Internet Registry

References

All RFCs are available from http://tools.ietf.org/html/

1. DARPA Internet Program. RFC791 - Internet Protocol Specification. Sep 1981
2. Internet Engineering Task Force. RFC1883 - Internet Protocol, Version 6 (IPv6) Specification. Dec 1995
3. Internet Engineering Task Force. RFC1918 - Address Allocation for Private Internets. Feb 1996

ABOUT THE AUTHOR

NIGEL TITLEY

Nigel Titley graduated from Oxford University in 1976 and went to work for British Telecom at Martlesham Heath in Suffolk. After leaving BT he moved on to work for various Internet service providers. He now works for Easynet. He helped form the London Internet Exchange (LINX) and Nominet and in 2003 he was elected onto the Board of the RIPE NCC.