

April 2013  
Geoff Huston

## **A Primer on IPv4, IPv6 and Transition**

There is something badly broken in today's Internet.

At first blush that may sound like a contradiction in terms, or perhaps a wild conjecture intended only to grab your attention to get you to read on. After all, the Internet is a modern day technical marvel. In just a couple of decades the Internet has not only transformed the global communications sector, but its reach has extended far further into our society, and it has fundamentally changed the way we do business, the nature of entertainment, the way we buy and sell, and even the structures of government and their engagement with citizens. In many ways the Internet has had a transformative effect on our society that is similar in scale and scope to that of the industrial revolution in the 19th century. How could it possibly be that this prodigious technology of the Internet is "badly broken?" Everything that worked yesterday is still working today isn't it? In this article I'd like to explain this situation in a little more detail and expose some cracks in the foundations of today's Internet.

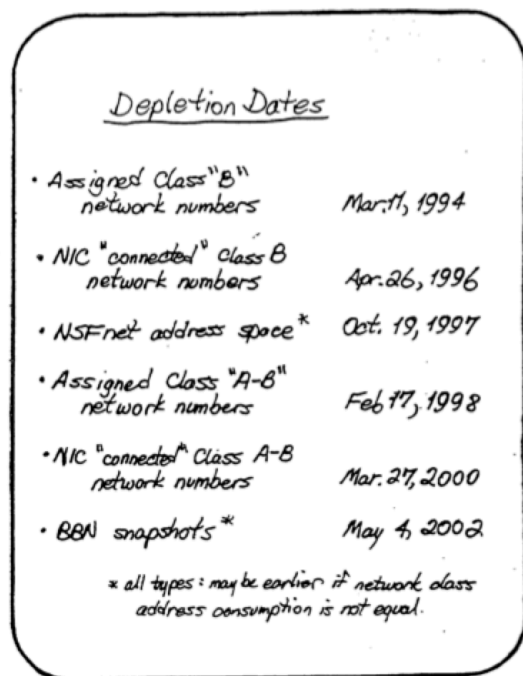
You see it's all about addresses. In a communications network that supports individual communications it's essential that every reachable destination has its own unique address. For the postal network it's commonly your street address. For the traditional telephone network it's your phone number. This address is not just how other users of the network can select you, and only you, as the intended recipient of their communication. It's how the network itself can ensure that the communication is correctly delivered to the intended recipient. The Internet also uses addresses. In fact the Internet uses two sets of addresses. One set of addresses is for you and I to use. Domain names are the addresses we enter into web browsers, or what we use on the right hand side of the @ in an email address. These addresses look a lot like words in natural languages, which is what makes them so easy for we humans to use. The other set of addresses are used by the network. Every packet that is passing through the Internet has a digital field in its header that describes the network address of the packet's intended delivery address: it's "destination address." This address is a 32 bit value. A 2 bit field has four possible values, a 3 bit field has eight possible values, and by the same arithmetic a 32 bit field has 2 to the power 32, or some 4,294,967,296 unique values.

If every reachable device on the Internet needs a unique address in order to receive packets, then does that mean that we can only connect at most some 4 billion devices to the Internet? Well, in general terms, yes! And once we reach that hard limit of the address size, should we expect to encounter problems? Well, in general terms, yes!

Running out of addresses in any communications network can pose a massive problem. We have encountered this a number of times in the telephone network, and each time we've managed to add more area codes, and within each area we've added more in-area digits to telephone numbers to accommodate an ever-growing population of connected telephone handsets. Every time we've made this change to the address plan of the telephone network we needed to reprogram the network. Luckily, we didn't need to reprogram the telephone handsets as well. We just had to re-educate telephone users to dial more digits. With care, with patience, and with enough money this on-the-fly expansion of the telephone system's address plan can be undertaken relatively smoothly. But this approach does not apply to the Internet. The address structure of the Internet is not only embedded into the devices that

operate the network itself, the very same address structure is embedded in every device that is attached to the network. So if, or more correctly, when, we run out of these 32 bit addresses on the Internet we are going to be faced with the massive endeavour of not only reprogramming every part of the network, but also reprogramming every single device that is attached to the network. Given that the Internet today spans more than 2.3 billion users and a comparable number of connected devices then this sounds like a formidable and extremely expensive undertaking.

If running out of IP addresses is such a problem for the Internet then you'd like to hope that we could predict when the ominous event would occur, and then give ourselves plenty of lead time to dream up something clever as a response. And indeed we did predict this address depletion. Some 23 years ago, in August 1990, when the Internet was still largely a research experiment and not the foundation bedrock of the global communications enterprise we saw the first prediction of address runout. At the time Frank Solensky a participant in the Internet Engineering Task Force (IETF) extrapolated the growth of the Internet from the emerging experience of the US National Science Foundation's NSFNET, and similar experiences in related academic and research projects, and predicted that the pool of addresses would run out in some 6-10 years time.



• Assigned Class "B" network numbers	Mar. 11, 1994
• NIC "connected" Class B network numbers	Apr. 26, 1996
• NSFnet address space*	Oct. 19, 1997
• Assigned Class "A-B" network numbers	Feb. 17, 1998
• NIC "connected" Class A-B network numbers	Mar. 27, 2000
• BBN snapshots*	May 4, 2002

\* all types: may be earlier if network class address consumption is not equal.

*Frank Solensky's Report on Address Depletion, Proceedings of IETF 18, p. 61, Vancouver, August 1990*  
(<http://www.ietf.org/proceedings/18.pdf>)

The technical community took this message to heart, and started working on the problem in the early 1990's.

From this effort emerged a stop gap measure that while it was not a long term solution, would buy us some urgently needed extra time. At the time the Internet's use of address use was extremely inefficient. In a similar manner to a telephone address that uses an area code followed by a local number part, the Internet's IP address plan divides an IP address into a network identifier and a local host identifier. At the time we were using an address plan that used fixed boundaries between the network identification part and the host identification part. This address plan was a variant of a "one size fits all" approach, where we had three sizes of host addresses within the network: one size was just too big for most networks, one size was too small, and the only one that left was capable of spanning an Internet of just 16,382 networks. It was this set of so-called "Class B" address blocks that Frank Solensky predicated to run out in four year's time.

So what was the stop gap measure? Easy. Remove the fixed boundaries in the address plan and provide networks with only as many addresses as they needed at the time. It was hoped that this measure would give us a few more years of leeway to allow us to develop a robust long term answer to this address problem. The new address plan was deployed on the Internet in early 1993, and for a couple of years it looked like we were precisely on track, and, as shown in Figure 2, this small change in the address plan, known as *Classless Inter-Domain Routing* (CIDR), would buy us around 2 or 3 years of additional time to work on a longer term approach to IP address exhaustion.




Figure 2 – CIDR and Address Consumption

As things turned out, we were wrong in that 2 – 3 year estimate.

The reason why we were wrong was that a second stop gap measure was also developed in the early 1990's. This new technology cut right to the heart of the architecture of the Internet and removed the strict requirement that every attached device needed its own unique address on the Internet.

The approach of *Network Address Translators* (NATs), allowed a collection of devices to share a single public IP address. The devices that were located “behind” a NAT could not be the a target of a new communication, so that, for example, you could not host a web service if you were behind a NAT, but as long as the devices behind the NAT initiated all communications, then the NAT function became invisible, and the fact that an IP address was being shared across multiple devices was effectively irrelevant. In a model of clients and servers, then as long as you only placed the clients behind a NAT then it was possible to share a single IP address across multiple clients simultaneously.

The emerging retail ISP industry took up this NAT technology with enthusiasm. The provisioning model for retail Internet services was for a single IP address provided for each connected service, which was then shared by all the computers in the home using a NAT that was embedded into the DSL or cable modem that interfaced the home network to the service provider network. The IP address consumption levels dropped dramatically, as it was no longer a case of requiring a new IP address for each connected device, but instead requiring a single IP address for each connected service. And as the home collected more connected devices, none of these devices drew additional addresses from the IP address pool.

Instead of buying a couple of years of additional breathing space to design a long term solution to address depletion, the result of the combination of classless addressing and NATs was that it looked like we had managed to push the issue of address depletion out by some decades! The most optimistic prediction of address longevity in around 2001 predicted that IPv4 address depletion might not occur for some decades, as the address consumption rate had flattened out, as shown in Figure 3.




Figure 3 – CIDR, NATs and Address Consumption

Perhaps it may have been an unwarranted over-reaction, but given this reprieve the industry appeared to put this entire issue of IP address depletion in the Internet onto the top shelf of the dusty cupboard down in the basement.

As events turned out, that level of complacency about the deferral of address depletion was misguided. The next major shift in the environment was the mobile Internet revolution of the last half of the 2000's. Before then mobile devices were generally just wireless telephones. But one major provider in Japan had chosen a different path, and NTT DOCOMO launched Internet-capable handsets onto an enthusiastic domestic market in the late 1990's. Their year-on-year rapid expansion of their mobile Internet service piqued the interest of many mobile service operators in other countries. And when Apple came out with a mobile device that included a relatively large well-designed screen and good battery life, an impressive collection of applications and of course a fully functional IP protocol engine, the situation changed dramatically. The iPhone was quickly followed by a number of other vendors, and mobile operators quickly embraced the possibilities of this new market for mobile Internet services. The dramatic uptake of these services implied an equally dramatic level of new demand for IP addresses to service these mobile IP deployments, and the picture for IP address depletion one more changed. What was thought to be comfortably far into the future problem of IP address depletion once more turned into a here and now problem.




Figure 4 – Address Consumption

Even so, we had exceeded our most optimistic expectations and instead of getting a couple of years of additional breathing space from these stop gap measures, we had managed to pull some 15 additional years of life out of the IPv4 address pool. But with the added pressures from the deployment of IP into the world's mobile networks we were once more facing the prospect of imminent address exhaustion in IPv4. So it was time to look at that long term solution. What was it again?

During the 1990's the technical community did not stop with these short term mitigations. They took the address depletion scenario seriously, and considered what could be done to define a packet-based network architecture that could span not just billions of connected devices but hundreds of billions of devices or more. Out of this effort came version 6 of the Internet Protocol, or IPv6. The changes to IPv4 were relatively conservative, apart from one major shift. The address fields in the IP packet header were expanded from 32 bits to 128 bits. Now every time you add a single bit you double the number of available addresses. This approach added 96 bits to the IP address plan. Yes, that's 340,282,366,920,938,463,463,374,607,431,768,211,456 possible addresses!

This approach to IPv6 appeared to adequately answer the need for a long term replacement protocol with enough addresses to fuel a rapacious silicon industry that can manufacture billions of processors each and every year. However, there was one residual annoying problem. The problem arises from one of the underlying features of the Internet's architecture: IP is an "end-to-end" protocol. There is no defined role for intermediaries in packet delivery. In the architecture of the Internet, what gets sent in a packet is what gets received at the other end. So if a device sends an IPv4 packet into the network, what comes out is an IPv4 packet, not an IPv6 packet. Similarly, if a device sends an IPv6 packet into the network then what comes out at the other end is still an IPv6 packet. The upshot of this is that IPv6 is not "backward compatible" with IPv4. In other words setting up a device to talk the "new" protocol means that it can only talk to other devices that also talk the same protocol. This device is completely isolated from the existing population of Internet users. What were these technology folk thinking in offering a new protocol that could not interoperate with the existing protocol?

What they were thinking was that this was an industry that was supposedly highly risk averse, and that once a long term replacement technology was available then the industry would commence broad adoption well before the crisis point of address exhaustion eventuated. The idea was that many years in advance of the predicted address exhaustion time, all new Internet devices would be configured to be capable of using both protocols, both IPv4 and IPv6. And the idea was that these bilingual devices would try to communicate using IPv6 first and fall back to IPv4 if they could not establish a connection in IPv6. The second part of the transition plan was to gradually convert the installed base of devices that only talked IPv4 and reprogram them to be bilingual in IPv6 and IPv4. Either that, or send these older IPv4-only devices to the silicon graveyard!

The transition plan was simple. The more devices on the Internet that were bilingual the more that the conversations across the network would use IPv6 in preference to IPv4. Over time IPv4 would essentially die out and support for this legacy protocol would be no longer required.

However one part of this plan was critical. We were meant to embark on this plan well before the time of address exhaustion, and, more critically, we were meant to complete this transition well before we used that last IPv4 address.




Figure 5 – The IPv6 Transition Plan

And to some extent this is what happened. Microsoft added IPv6 to its operating systems from the mid 2000's with the Windows Vista and Windows Server 2008 products. Apple similarly added IPv6 into their Mac OSX system from around 2006. More recently, IPv6 support has been added into many mobile devices. These days it appears that around one half of all devices connected to the Internet are bi-lingual with IPv6 and IPv4. This is indeed a monumental achievement, and much of the effort in re-programming the devices that are attached to the Internet to speak the new protocol has been achieved. So we are all ready to switch over the Internet to use IPv6, yes? Well, no, not at all.

So what's gone wrong?

Many things have not gone according to this plan, but perhaps there are two aspects of the situation that deserve highlighting here.

Firstly, despite the addition of IPv6 into the popular computer platforms, the uptake of IPv6 in the network is just not happening. While there was a general view that the initial phase of IPv6 adoption would be slow, the expectation was that the use of IPv6 would accelerate along exponentially increasing lines. But so far this has not been all that evident. There are many metrics of the adoption of IPv6 in the Internet, but one of the more relevant and useful measurements is that relating to client behaviour. When presented with a service that is available in both IPv4 and IPv6, what proportion of clients will prefer to use IPv6? Google provide one measurement point, that measures a sample of the clients who connect to Google's service. Their results are shown in Figure 6.




Figure 6 – IPv6 Adoption (<http://www.google.com/intl/en/ipv6/statistics/>)



Over the past four years Google has seen this number rise from less than 1% of users in early 2009 to a current value of 1.2%. Its one of those glass half-full or half-empty stories. Although in this case it's the glass is either 1% full or 99% empty! If broad scale use of IPv6 is the plan, then right now we seem to be well short of that target. On a country-by country basis the picture is even more challenging. Only 9 countries have seen the proportion of IPv6 users rise above 1%, and the list has some surprising entries.

IPv6 Users by Country					
Date: 18 Apr 2013					
ISO-3166 Code	Internet Users	V6 Use ratio	V6 Users (Est)	Population	Country
RO	8656148	10.02%	867346	22082012	Romania
LU	469505	6.01%	28217	513682	Luxembourg
FR	50185773	5.81%	2915793	65007479	France
JP	100762512	3.31%	3335239	125953141	Japan
DE	67933822	2.60%	1766279	82144888	Germany
US	249503375	2.44%	6087882	318650543	United States of America
CZ	7210754	2.18%	157194	10170316	Czech Republic
PE	10537762	1.40%	147528	30902529	Peru
BE	8503710	1.31%	111398	10446819	Belgium

Figure 7 – IPv6 Adoption (<http://labs.apnic.net/dists/v6dcc.html>)

Its hard to portray this as evidence of broad based adoption of IPv6. Its perhaps more accurate to observe that a small number of network providers have been very active in deploying IPv6 to their customer base, but these providers are the minority, and most of the Internet remains locked deeply in IPv4. If a significant proportion of the end devices support IPv6 then why are these use metrics so unbelievably small? It appears that the other part of the larger network re-programming effort, that of enabling the devices sitting within the network to be IPv6-capable, has not taken place to any significant extent. It's still the case that a very large number of ISPs do not include IPv6 as part of their service offering, which means that even if an attached computer or mobile device is perfectly capable of speaking IPv6, if the access service does not support IPv6 service then there is effectively no usable way for the device to use IPv6. And even when the service provider supplies IPv6 as part of its service bundle, it may still be the case that the user's own network devices, such as the in-home NAT/modems and other consumer equipment that supports in in-home networks, such as a WiFi base station or a home router may only support IPv4. Until this equipment is replaced or upgraded, then IPv6 cannot happen. The result is as we seen in the IPv6 usage metrics today: when offered a choice between IPv4 and IPv6, some 99% of the Internet's connected devices will only use IPv4.

Secondly, we've now crossed into a space that was previously regarded as the unthinkable: we've started to run out of IPv4 addresses in the operating network. This address exhaustion started with the central address pool, managed by the Internet Assigned Numbers Authority (IANA). The IANA handed out its last address block in February 2011. IANA hands out large blocks of addresses (16,777,216 addresses per "block") to the Regional Internet Address Registries (RIRs), and in February 2011 it handed out the last round of address blocks to the RIRs. Each of the five RIRs operates independently, and each will themselves exhaust their remaining pool of IPv4 addresses in response to regional demand. APNIC, the RIR serving the Asia Pacific region, was the first to run out of addresses, and in mid April 2011 APNIC handed out its last block of "general use" IPv4 addresses. (as a side remark here, APNIC still had 17 million addresses held aside at that point, but the conditions associated with allocations from this so-called "final /8" are than each recipient can receive at most an allocation of a total of just 1,024 addresses from this block.) This represented an abrupt change in the region. In the last full year of general use address allocations, 2010, APNIC consumed some 120 million addresses. In 2012, the first full year of operation under this last /8 policy the total number of addresses handed out

in the region dropped to 1 million addresses. The unmet address demand from this region appears to be growing at a rate of around 120 – 150 millions addresses per year.

The region of Europe and the Middle East has been the next to run out, and in September 2012 the RIPE NCC, the RIR serving this region, also reached its “last /8” threshold, and ceased to hand out any further general use IPv4 addresses. The process of exhaustion continues, and the registry that serves Northern America and parts of the Caribbean, ARIN, has some 40 million addresses left in its address pool. At the current consumption rate ARIN will be down to its last /8 block 12 months from now, in April 2014. LACNIC, the regional registry serving Latin America and the Caribbean, currently has some 43 million addresses in its pool, and is projected to reach their last /8 slightly later in August 2014. The African regional registry, AFRINIC, has 62 million addresses, and at its current address consumption rate, the registry will be able to service address requests for the coming seven years.




Figure 8 – IPv4 Address Depletion (<http://labs.apnic.net/ipv4/report.html>)

So if the concept was that we would not only commence, but complete the process of transition to use IPv6 across the entire Internet before we got to that last IPv4 address, then for Europe, the Middle East, Asia and the Pacific this is not going to happen. It's just too late. And for North and South America it's also highly unlikely to happen in time.

And the slow pace of uptake of IPv6 points to the expectation that this “running on empty” condition for the Internet address plan may well continue for some years to come.

We are now entering into a period of potential damage for the Internet. If the objective of this transition from IPv4 to IPv6 was to avoid some of the worse pitfalls of exhaustion of the IPv4 address space in the internet, then we've failed.

The consequence of this failure is that we are now adding a new challenge for the Internet. It's already a given that we are meant to sustain continued, and indeed accelerating, growth in terms of the overall size of the network and the population of connected devices. The pace of this growth is expressed as a demand for some 300 million additional IP addresses per year, and the figures from the device manufacturers point to a larger figure of some 500 – 700 million new devices being connected to the Internet each year. And the number grows each year. We are expanding the Internet at ever faster rates. As if riding this phenomenal rate of growth on the existing infrastructure and existing technology base wasn't challenging enough, we also have the objective not just to maintain, but to accelerate the pace of transition to IPv6. These two tasks were already proving to be extremely challenging, and we've been slipping on the second. But we now have the additional challenge of trying to achieve these two objectives without the supply of any further IPv4 addresses. At this point the degree of difficulty starts to get uncomfortably close to ten!



This situation poses some architectural consequences for the Internet. Until now we've managed to push NATs out to the edge of the network, and make address compression something that end users did in their home networks. The consequences of failure of such devices and functions are limited to the edge network served by the NAT. We are now deploying mechanisms that allow this NAT function to be performed in the core of the carriage networks. This introduces a new set of unquantified factors. We've little experience in working with large scale NAT devices. We have no idea of the failure modes, or even the set of vulnerabilities in such an approach. We are still debating the appropriate technical approach in the standards bodies, so there are a variety of these service provider NAT approaches being deployed. Each NAT approach has different operational properties, and different security aspects. But now we don't have the luxury of being able to buy more time to explore the various approaches and understand the relative strengths and weaknesses of each. The exigencies of address exhaustion mean that the need for carrier level NAT solutions is now pressing, and given that this is a situation that we never intended to experience, we find ourselves ill-prepared to deal with the side effects from this subtle change in the network's architecture. The greater the level of complexity we add into the network, and the wider the variation in potential network behaviours as a result, the greater the burden we then place on applications. If the network becomes complex to negotiate then applications are forced to explore the local properties of the network environment in order to provide the user with a robust service.

If the hallmark of the Internet was one of efficiency and flexibility based on a simple network architecture, then as we add complexity into the network what we lose is this same efficiency and flexibility that made the Internet so seductively attractive in the first place. The result is a network that is baroquely ornamented, and one that behaves in ways that are increasingly capricious.

We are hopelessly addicted to using a network protocol that has now run out of addresses. At this point the future of the Internet, with its projections of trillions of dollars of value, with its projections of billions of connected silicon devices, with its projections of petabytes of traffic, with its projections of ubiquitous fibre optics conduits spanning the entire world is now entering a period of extreme uncertainty and confusion. A well planned path of evolution to a new protocol that could comfortably address these potential futures is no longer being followed. The underlying address infrastructure of the network is now driven by scarcity rather than abundance, and this is having profound implications on the direction of evolution of the Internet.

There really is something badly broken in today's Internet.

---

## Disclaimer

The views expressed are the author's and not those of APNIC, unless APNIC is specifically identified as the author of the communication. APNIC will not be legally responsible in contract, tort or otherwise for any statement made in this publication.

---

## Author

*Geoff Huston* B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the

Australian academic and research sector. He is author of a number of Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005, and served on the Board of Trustees of the Internet Society from 1992 until 2001.

*[www.potaroo.net](http://www.potaroo.net)*