November 2013
Geoff Huston

IPv6 at the OECD

A Public Policy Perspective on IPv6

> Over the past few years I have attended a number of meetings of the OECD's Working Party on Communication Infrastructures and Services Policy as a member of the Internet Technical Advisory Committee. Their meetings have covered many aspects of communication infrastructure from spectrum management issues to "smart" televisions, and, unsurprisingly, the topic of the exhaustion of IPv4 addresses and the transition to IPv6 has been part of this Working Party's agenda.

The Organisation for Economic Co-operation and Development, the OECD, is a widely referenced and respected source of objective economic data and comparative studies of national economies and economic performance. The organization has a very impressive track record of high quality research and a justified reputation of excellence in its publications, even with its evident preference for advocating economic reform through open markets and their associated competitive rigors. OECD activities in the past have proved to be instrumental in facilitating change in governmental approaches to common issues that have broad economic and social dimensions.

Within the Directorate of Science Technology and Industry, the Internet has proved to be a fertile are of study. In the past the OECD has examined a broad spectrum of issues under the umbrella topic of the digital economy, including broadband bundling, big data, data-driven innovation, information security and privacy.

> The collection of published OECD papers that address various aspects of research into the digital economy can be found at http://www.oecd-ilibrary.org/science-and-technology/oecd-digital-economy-papers_20716826

So how does IPv6 fit into this picture of OECD activities?

The origins of IPv6 were based in work initiated by the technical community, and could be described as a technical response to an anticipation of future need, rather than as a solution to a current need. More than two decades ago, just as the prospects for success for the Internet were firming up, the prospects for its eventual demise were also becoming apparent. While the Internet Protocol was capable of embracing a network populated with more than 4 billion devices, real network deployments ware incapable of achieving such perfection of address utilization efficiency. What was deployed was a network that was effectively capable of serving at most some hundreds of millions of connected devices. At the same time computing devices continued along a path of getting smaller and cheaper, so the prospect of a world populated with billions of chattering devices was conceivable, and we would need some form of network technology that was capable of scaling to this same relative size. The

Internet Engineering Task Force (IETF) decided in the early 1990's to devise a new version of the Internet Protocol that would scale across a truly massive population of connected devices that readily accommodated a network spanning hundreds of billions of connected devices. At the same time the IETF embarked on developing some short term technical mitigations, intended to decrease the rate of consumption of protocol address, which would buy some time while this new protocol was being developed. The idea at the time was that we would never actually run out of network protocol addresses — we would've completed the migration of the Internet to use this revised protocol, called IPv6, well before we got to use that last address in the existing protocol, IPv4.

Two decades later its now apparent that we've deviated from this plan. The regions of Asia, Oceania, Europe and the Middle East have now effectively exhausted their supply of IPv4 addresses, and are now attempting to service an ever expanding Internet without any further IPv4 addresses to support this ongoing growth in demand. We anticipate that North and South America will be in a similar position of address exhaustion by the end of 2014.

Why have we managed to run out of IPv4 addresses? Why haven't we switched over to use IPv6 by now?

Given that this was a technical response to the perceived problem, should we ask the question as to whether there is some residual failure in the technical process of developing this new Internet protocol? Is there something in IPv6 that prevents its broadscale deployment?

IPv6 is well specified technically. The changes introduced by this protocol are limited in scope to the internetworking layer of the overall Internet Protocol architecture, and the upper levels o the protocol stack, principally TCP and UDP, are unaltered. Applications that operate in IPv4 will, in general, operate in IPv6. Vendors can and do provide equipment that support both IPv4 and IPv6 protocols. The major vendors of operating systems for various conventional computing platforms all now support IPv6, as do many mobile devices. All of the elements appear to be in place, and in many way there has been a better orchestration of supply channels for this new protocol than was the case 20 years ago with the rise of the IPv4 network infrastructure. So the issues at play here do not appear to be specifically issues related to some failure in the technology base of the network.

Can we find some answers in the business and policy environment of the Internet?

The factors that have lead us into this anomalous situation can be considered from a business perspective within the context of a largely deregulated and highly competitive market environment. It is often the case that the successful adoption of new technologies have been accompanied by some form of early adopter advantage. The higher risks associated with using a new technology are offset against some form of advantage to be gained by early incumbency in an emerging market for those goods and services that are derived from this technology, assuming that the technology achieves widespread adoption. But IPv6 has little to distinguish itself from IPv4 in terms of services and service costs, so the incentives for early adoption provided insufficient grounds for many providers to make the decision to commence widespread deployment. The major incentive left to propel deployment of IPv6 was that of "future risk", where the prospect of exhaustion of the continued supply of IPv4 addresses would be associated with an escalating cost to providers to deploy various forms of address sharing technologies to allow their finite pools of IPv4 address to be shared across an expanding pool of customer devices. However, this has not yet proved to be a major factor in the incentives to drive further deployment of IPv6 in the Internet.

There is one further factor to consider here, namely that of backward compatibility. An IPv4-speaking host cannot address IP packets directly to an IPv6-speaking host, and the same in reverse in that an IPv6-speaking host cannot address packets directly an IPv4-speaking host. The major implication from this lack of backward compatibility is that it not feasible to switch over a network from IPv4 to IPv6. Instead, the transition involves a "dual-stack" phase, where IPv6 is turned on in parallel with continued operation of IPv4. Networks need to continue to support IPv4 for as long as there are significant levels of IPv4-only networks and services still deployed. Once the overall majority of the Internet supports a dual-stack environment can networks start to turn off their continued support for IPv4. Therefore, while there is no particular competitive advantage to be gained by early adoption of IPv6, the timing of the ultimate end point, namely the decommissioning of IPv4, is not determined by the early adopters, but, conversely, by the late adopters. It would appear, at least superficially, that if dual stack operation incurs some additional cost, then its possible to make the case that late adopters may be advantaged by this scenario, as they would need to support a dual stack environment for the shortest time period. On the other hand, if such actors leave it too late, then the risk is that their customers would be isolated should the mainstream set of networks commence dropping IPv4 support in the final phases of their transition to IPv6, so late adoption is not a strategy without some risk. Whether this late adoption is a major contributory factor or not, the observation still remains that we have not managed to avoid the situation of attempting to support an expanding network that had exhausted its supply of further IPv4 end point addresses, and the Internet is now running on empty.

However this observation does not seem to have had any visible repercussions. The Internet still looks much the same for most of its 2 billion-strong user base. Somehow we've managed to preserve the functionality of applications and services across the Internet while not having enough unique end point addresses to go round. To achieve this somewhat remarkable feat we've relied heavily on the original short term mitigations, originally devised in the early 1990's, and have transformed them into integral components of today's internet infrastructure. Much of the recent growth of the Internet in the world of mobile services has been supported using Carrier Grade Network Address Translators (CGNs), where a pool of customer connections are configured to share a far smaller pool of IPv4 addresses. While such NAT devices have been used at the periphery of the network for more than a decade, we have taken this approach and applied it to the access providers' network. This approach has been the conventional approach for many of the mobile 3G networks since their initial deployment, and it is now uncommon to see wireless networks that present public IPv4 addresses to their connected devices. With the increasing pressures of continued growth of the Internet, and the continued deferral of widespread adoption of IPv6, we are now seeing a number of operators contemplating widespread deployment of CGNs in both the wired as well as the wireless network service domains.

There are some serious concerns about the long term implications of an Internet that increases its reliance on CGNs in this manner as a response to an address-depleted Internet. A recent study from OFCOM in the UK, "Report on the Implications of Carrier Grade Network Address Translators", (http://stakeholders.ofcom.org.uk/binaries/research/technology-research/2013/cgnat.pdf) points to risks of an emerging picture of incumbent dominance and decreasing competitive pressure in the Internet Service Provider sector, and, perhaps more seriously, an emerging ability for access carriage providers to decide what applications and services that will be accessed by their users, and the consequent ability for the access sector to impose terms and conditions on content providers in order for these content providers to reach their end customer base. This is coupled with considerations of a potential shift in position of access providers from that of common carriers to one of assumed liability for the applications and services that are passed through their CGNs.

Were we to actually experience these outcomes, it would pose some serious questions for public regulators and policy makers. The risks of a failure in the ability of the internet to maintain its essential openness and decentralized availability, which would be a possible consequence of a failure to complete a timely transition to IPv6 across the entirety of the Internet, could lead to the contemplation of the prospect of a market failure in the larger Internet economy. The risks inherent in such a scenario underline the importance of the commitment made by Ministers in the Seoul Declaration on the Future of the Internet Economy (http://www.oecd.org/sti/40839436.pdf), and the need to: "Encourage the adoption of the new version of the Internet protocol (IPv6), in particular through its timely adoption by governments as well as large private sector users of IPv4 addresses, in view of the ongoing IPv4 depletion."

But if the risks of failure are so dire then why stop at "encouraging" the adoption of IPv6? Why shouldn't the national regulator direct its local Internet sector to adopt IPv6 through administrative fiat?

Part of the sensitivity here is that the time when the public telecommunications environment was owned and operated entirely within the public sector has long since disappeared. The deregulation of this sector has allowed not only the Internet to thrive through the significant injection of private capital, but we've also seen a rapid transformation of many economies in response to the opportunities opened up by the Internet. Hundreds of thousands of small businesses are now international exporters thanks to eBay and similar online market facilitators. Entrepreneurial activities are flourishing with the adoption of crowd sourcing models of venture capital. But of course within all this private sector activity, the ability of the public sector to regulate and direct such activities has also been also altered significantly. Much of the direction of the Internet is now being undertaken though the interaction of competitive pressures and market forces, and the ever present risks of aggregation and market dominance are being countered by constant innovative pressure from new market entrants. But this relies on these same market forces preserving an environment that remains open to innovation and competition. The challenge for the public regulator in this sector is to steer a careful course between allowing the sector to operate efficiently and facilitate further private sector investment in the provision of infrastructure and services that support national economic activity, while at the same time ensuring that the public interest is preserved through the efficient provision of goods and services. The days of the public sector "directing" activities in the public communications sector are over in many national economies, and these days it's more common to see language related to "encouraging" and "facilitating" take its place.

Within the context of the studies undertaken by the OECD members of the OECD's Internet Technical Advisory Committee (ITAC) (http://www.internetac.org) have been active within the OECD's Working Party on Communication Infrastructures and Services Policy (CISP) in highlighting the risks associated with a failure to complete a timely transition of the Internet to IPv6. ITAC members have prepared a study for the Working Party that describes in some detail the current state of the Internet's transition to IPv6, and provides an analysis on this situation. ITAC members have also taken the lead in drafting recommendations to the OECD on possible actions by OECD Member Countries, and others, that would hasten the deployment of IPv6 through various potential public sector initiatives. Such measures may not necessarily propel accelerated IPv6 deployment in and of themselves, but the public sector is a significant consumer and beneficiary of online services, and recommending that this sector make determined steps to use IPv6 increases the likelihood of success of the overall process of transition to IPv6.

Disclaimer

The views expressed are the authors' 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.

About the 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