

OECD (2014-04-03), "The Internet in Transition: The State of the Transition to IPv6 in Today's Internet and Measures to Support the Continued Use of IPv4", *OECD Digital Economy Papers*, No. 234, OECD Publishing, Paris.
<http://dx.doi.org/10.1787/5jz5sq5d7cq2-en>



OECD Digital Economy Papers No. 234

The Internet in Transition: The State of the Transition to IPv6 in Today's Internet and Measures to Support the Continued Use of IPv4

OECD

FOREWORD

This report was presented to the OECD Working Party on Communication, Infrastructures and Services Policy (CISP) in June 2013. The Committee for Information, Computer and Communications Policy (ICCP) approved this report in December 2013 and recommended that it be made available to the general public. It was prepared by Geoff Huston, Chief Scientist at the Asia Pacific Network Information Centre (APNIC). The report is published on the responsibility of the Secretary-General of the OECD.

Note to Delegations:

*This document is also available on OLIS under reference code:
DSTI/ICCP/CISP(2012)8/FINAL*

TABLE OF CONTENTS

FOREWORD.....	2
THE INTERNET IN TRANSITION: THE STATE OF THE TRANSITION TO IPV6 IN TODAY'S INTERNET AND MEASURES TO SUPPORT THE CONTINUED USE OF IPV4.....	4
Abstract.....	4
MAIN POINTS	5
INTRODUCTION.....	8
Address management and IPv4 runout	11
Addresses and address structures.....	11
Description of the address distribution framework.....	12
Statistics on address distribution.....	12
"Used" and "unused" addresses	13
Predictions on IPv4 runout.....	14
Network address translation as a response to IPv4 address exhaustion.....	17
The IPv6 protocol	20
(a) What changed with IPv6.....	20
(b) And what did not	20
(c) The demand drivers for the IPv6 transition	21
(d) Backwards compatibility issues.....	22
The IPv6 transition plan.....	22
(e) Assumptions	23
(f) Dual-stack and IPv4 address exhaustion	24
Measuring the state of the IPv6 transition	25
Measurement using the routing system.....	25
(g) Measurement using the domain name system	27
Measurement using Internet Traffic Statistics	28
Measurements of end client capabilities	29
Potential scenarios	31
(h) Openness and innovation.....	32
(i) Carrier grade NATs	32
Conclusion	33
GLOSSARY OF TERMS.....	35
NOTES	37

THE INTERNET IN TRANSITION: THE STATE OF THE TRANSITION TO IPV6 IN TODAY'S INTERNET AND MEASURES TO SUPPORT THE CONTINUED USE OF IPV4

Abstract

This report considers the transition from IPv4 to IPv6 alongside the use of network technologies to prolong IPv4 use in the face of depletion of further IPv4 protocol addresses. The report neither aims to address all the issues surrounding the transition to IPv6 nor to detail the economic incentives faced by various Internet actors. The report first provides a status update of address management issues and the run-out of IPv4. It then describes the advantages and limitations of increased use of network address translation as one response to sustain the use of IPv4 in the face of IPv4 address exhaustion. The following sections provide an overview of the IPv6 protocol; the advantages of IPv6 deployment as a response to IPv4 address exhaustion and the IPv6 transition plan compared to actual deployment to date. Finally, the report examines the choices facing individual actors, their potential consequences, and the policy implications on openness and innovation for the future of the Internet.

MAIN POINTS

The Internet can be regarded as a classic case of an experiment undertaken in a research laboratory turning into a runaway public success. By virtue of an open, edge-based decentralised architecture, the Internet has experienced a rapid trajectory of adoption, such that it has now assumed the role of the common digital communications technology for economies and societies across the world.

This particular version of the Internet communications protocol was not designed for this particular all-embracing role. As early as 1990, concerns were raised in the technical community that certain aspects of this technology would not readily scale into ubiquitous global deployment in a world that contained dense concentrations of communicating devices. In particular it was felt that the scope of 4 billion addresses defined within this communications protocol was simply insufficient for such a role. The technical response to these concerns was to define a new version of the Internet Protocol that would encompass a far larger address domain, namely version 6 of the Internet Protocol, or IPv6. The technical specification of IPv6 was largely completed by 1996. At this time the technical community initiated a programme to promote the need to transition the Internet to use IPv6, and proposed a deadline for completion of this transition well in advance of the anticipated exhaustion of the IPv4 address pool. The transition process had one major challenge, in that IPv6 is not backward compatible with IPv4. This implies that this transition is not a simple case of replacing IPv4 with IPv6 and moving on. Networks, edge devices, and service delivery points all need to support operation in both protocols simultaneously, and do so until the entire Internet has this dual protocol capability. Only then can the IPv4 component of the network be phased out of service.

As this report illustrates with various measurements, the transition to IPv6 is still in its early phase. Meanwhile, the Internet has continued with accelerating growth, as each year sees greater levels of deployment of connected devices than previous years. The initial wave of expansion of the wired access network has been complemented in recent years by the co-opting of mobile telephony services by the hand-held Internet device. This expansion of the scope of the network places further pressures on the dwindling pool of available IPv4 addresses, to such an extent that the ongoing supply of IPv4 addresses that has fuelled the continued growth of the Internet has now ceased in Europe, the Middle East, Asia and Oceania. At the same time the transition to IPv6, through a dual-stack capable Internet, has not followed quite the same path of widespread rapid deployment.

Some parts of the industry have appreciated the nature of the issues relating to running out of addresses within the network's infrastructure, and have planned and acted accordingly. Microsoft's widespread Windows products have included an IPv6 protocol engine alongside their IPv4 engine in their operating system since late 2001. Similar moves have been made by Apple in their MAC OSX operating system, and support is also to be found in Linux systems, and Linux-based derivative platforms, such as Android. Much of the infrastructure of the network, including backbone transmission paths and the Domain Name System (DNS) infrastructure, is also IPv6-capable. Not all parts of the industry, however, have interpreted this message about the need for transition to IPv6 as an imperative call for action. The result is that while it appears that over one half of the end-user equipment deployed on the wired Internet is capable of supporting IPv6 today, less than 2% of this same equipment is able to support connections to external services using IPv6.¹

Similarly, a number of providers have made substantial efforts to extend IPv6 services through their access networks. Notable efforts were made by a number of access providers, including RDSnet in Romania, KDDI in Japan, Free.fr in France, Comcast, Verizon Wireless, Savvis, T-Mobile and AT&T in the United States, Deutsche Telekom and Kabel Deutschland in Germany and XS4ALL in the Netherlands.

All these providers have recently made significant efforts to extend IPv6 access to consumers through their mass-market Internet access services. However, IPv4 address depletion has imposed a further agenda on many service providers, many of whom now appear to be concentrating their current efforts on the deployment of technologies that will conserve their remaining IPv4 address stocks by deploying address-sharing middleware that will permit the sharing of IPv4 addresses across multiple customers. This potential change in the immediate agenda for many entities in this sector may have profound implications for the future of the Internet.

The Internet's position as a critical component of a novel and valuable information economy has assumed that this area will be subject to continued evolution; as such decentralised open platforms are inevitably the constant focus of innovation. Such innovation not only creates greater efficiencies in Internet-based services, but also reaches out to an ever broader set of activities and pastimes. This form of innovation is decentralised and opportunistic, as well as being entirely open. With an open end-to-end network, any form of communications can be conceived and implemented. Competition in the provision of goods and services in this environment means not only realising greater efficiencies in the existing processes of service provision, but also enables production of the equivalent services using entirely novel approaches.

The degree of flexibility that is achievable in an open – end-to-end - network architecture is not readily sustainable in a networked environment where parts of the network use different communications protocols. The protocol translation functions at the interface between such different protocol realms generally impose constraints, in so far as the translation function tend to produce outcomes that are related to what is common across the protocols, while clients in either protocol realm cannot access the full functions and services available in the realm of the other protocol. Today's Internet exhibits some of these limitations, in that large parts of the Internet are operating in an environment of IPv4-only, while other sectors are embarking on transitioning their networks to IPv6. In many parts of the world, although with a significant number of notable exceptions, consumer market Internet service providers are not, as yet, deploying IPv6 infrastructure in their mainstream product set, and thereby not delivering IPv6 capability to end users. Instead, many actors in this sector are deploying IPv4 address extension technologies to share now scarce IPv4 addresses across multiple customers and modify their networks accordingly.

If this impediment to the momentum of IPv6 transition continues, it poses some long-term risks to the Internet and the global Internet economy. The dense proliferation of so-called middleware impairs the open clarity of the end-to-end architecture, and impedes many forms of innovation in the network service model. Damage to the ability of the Internet to remain open and receptive to further innovation must inevitably constrain future forms of competition, and increase the risk of market dominance by established carriage and content service providers.

The regulatory issues that would be generated by this particular scenario, were it to eventuate, could be relatively serious. If the Internet were no longer able to support the continued entrance of new competitors, and were it to be unable to allow such new entrants to innovate on the basic themes of service delivery and customer fulfilment, then policy makers and regulators may be faced with a scenario of growing market dominance of the means of access to digital services and resources, and their circulation by a small number of actors.

However, it should be noted that such a scenario is by no means the only possible outcome from the current situation. It is also possible that the existing efforts to maintain an open and flexible Internet, chiefly through the adoption of IPv6, continue and gather further momentum across the Internet. For IPv6 in particular this would imply that the increasing volumes of IPv6 deployment would generate economies of scale that would add further impetus for other actors to embark on a similar course of action. This does not in itself ameliorate the larger dimension of the risks involved in a failure to complete a timely transition

to IPv6, but this consideration of a gathering momentum of IPv6 adoption within the service provider sector increases the likelihood of the Internet continuing as an open, accessible, competitive and innovative platform.

The risks of a technical failure in the ability of the Internet to maintain its essential openness and decentralised 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, 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."²

INTRODUCTION

For more than two decades now almost every plot of some metric related to the Internet has a similar shape, namely the same curve rising steeply to the right over time. Across all of these graphs the overall story is consistent in that the Internet is still in a period of rapid growth and there are no visible signs of a slowing down of this expansion.

How big can the Internet grow? Is it like a theory of an ever-expanding cosmos, expanding for all eternity, or are there fundamental limits to the technology of the Internet, which will inevitably cause this growth to slow down.

When the Internet engineering community first studied this issue some two decades ago the conclusion was that the most pressing problem of potential exhaustion was that of the IP address space. While 32 bits of address space can mathematically encompass some 4.3 billion uniquely addressed devices, in the real world address space cannot be used as efficiently, and a more realistic limit to the number of addressable devices in a 32 bit Internet address space is somewhere between 40 million and 1 billion devices. At the time of the initial study the sector was using addresses relatively inefficiently, and the lower 40 million limit looked like it would be reached sometime in the mid to late 1990's.³

A number of recommendations from this study were adopted by Internet Service Providers (ISPs), including in 1994 the use of classless inter-domain addressing (CIDR), a technique intended to improve the address utilisation efficiency in the Internet by making its routing system more flexible (fuller, 2006).⁴ The major outcome of this effort occurred in late 1995 with the publication of the initial specification of a new version of the Internet protocol, IP version 6 (IPv6).⁵ The major change in this version of the IP protocol was the expansion of the IP address fields in the header from 32 bits per address to 128 bits. There is no doubt that 128 bits is a very large number, and if there was an address space shortfall in IP version 4 (IPv4), 128 bits of address is certainly a very large compensatory measure.

The expectation at the time was that ISPs would react prudently to news of the impending exhaustion of the IPv4 address pool and would collectively switch over to an all-IPv6 network well before the remaining pool of available IPv4 addresses was depleted. Unfortunately, the technical and economic incentives-related challenges associated with converting large numbers of interconnected IP networks to IPv6 were not adequately appreciated. As a result, this has not been the case. The central IPv4 address pool, managed by the operator of the Internet Assigned Numbers Authority (IANA) was fully allocated in February 2011, and, by August 2013, two of the five Regional Internet Registries had reached critical low points in their locally managed IPv4 address pools. The Internet, however, continues to rely on IPv4 for the overall majority of end users and services, and the visible levels of IPv6 deployment remain extremely low.

Why has the available IPv4 address pool been depleted to this level without the broad scale adoption of a viable alternative protocol, in the form of IPv6? An answer to this question lies, in part, in the de-coupled nature of the supply chain for the Internet's service provider sector. Integration of robust support for IPv6 into all forms of service provider, enterprise and customer equipment has taken many years. One factor for this is a lack of clear and coherent signalling within the supply chain. For example, many service providers prudently delayed their plans for IPv6 deployment in their service platforms until they could assure themselves that their equipment could support IPv6 at scale without any negative impact on their service's performance and robustness. The equipment vendors noted, in turn, that their developmental efforts were focussed on aspects of their equipment that influenced buyer's purchasing decision. Unless their customers made support for IPv6 a critical component of their purchasing decisions,

then equipment vendors would not necessarily place robust support for IPv6 uppermost in their product development schedules.

Another part of the explanation of the current situation lies in another outcome of the work in the early 1990's on address exhaustion, namely that on Network Address Translation (NAT).⁶ NAT units conventionally sit at the edge of the network and allow multiple interior end devices to share a single exterior public address. They act in a manner that allows outbound connections to be made by interior devices, and allows the subsequent bidirectional traffic flow associated with the connection to be supported, but they conventionally do not permit inbound connections to be initiated by exterior devices. This behaviour has been conventionally seen as a measure that assists in securing end systems from various forms of address-scanning malware. For ISPs this has become a very common model for Internet service deployment on wired networks. ISPs provide the user's edge device (such as, for example, a cable modem, or more generally any form of Customer Premises Equipment) with a single IPv4 address, and the user's device shares this single address across multiple devices in the user's local network by virtue of an inbuilt NAT. Doing so externalised the costs of IPv4 address scarcity to the user, and the shortcomings of the NAT model were mitigated by the perception of increased security of the user's network. The other advantage of this approach is that of incremental deployment, where each private network can make a local decision to use private address space rather than public IP address space, and use a NAT interface to the Internet, where the effect of this decision for use a NAT is also purely local to that private address realm. Such piecemeal local decisions to use NATs can be made far more readily than co-ordinating the process of making a global decision to change the underlying protocol used by everyone.

Much of the Internet's device population is now located behind NATs located at the edge of the network. To further increase the efficiency of use of a public IPv4 address, one approach is to move from a model of sharing a single address across the devices used by a single service provider's customer to a model of sharing an IPv4 address across the devices used by multiple customers. This is a "cascading NAT" model, where the edge NATs already in place are complemented by a NAT placed within the ISP's infrastructure. This second NAT, that enables address sharing across discrete customer endpoints, is commonly referred to as a "Carrier Grade NAT" or CGN. Again, this can be deployed incrementally as a local decision, as neither the edge NATs that sit behind the CGN, nor any other element of the network is directly impacted by this decision to deploy a CGN within the service provider's network. CGNs are not necessarily restricted to those NATs that complement existing edge-based NATs. Other forms of CGNs effectively replace the edge-based NAT, pooling the NAT behaviours at the edge to an aggregate NAT within the service provider network (such as is used in the Dual Stack-Lite IPv6 transition architecture). The common property of CGNs is that the NAT function is one that is provided within the service provider's infrastructure.

This raises the question as to whether CGNs represent a durable and fully functional response to the overall scaling problem posed by inexorable growth of the Internet. The increasing deployment of edge NATs and CGNs, with the emergence of networks with distinct address realms causes changes in the Internet's architecture that, in turn, generates changes in the service model for the Internet.

Widespread use of NATs cannot support a rich functional communications model. If it is believed that the potential value of computer-mediated communications Internet is more than that of simple client server transaction model of today's web, then why is IPv6 not being used to a greater extent already? What is holding back the Internet from completing the transition to a protocol that solves the address scaling issue, provides functional diversity in service models, and offers a larger set of security mechanisms that are missing with the intensive use of NATs with IPv4? Part of the issues that holding the Internet back is one of the nature of the transition itself. Unfortunately, IPv6 is not a backwards compatible variant of IPv4, and the two protocols do not interoperate directly. IPv6 is not a direct substitute for IPv4, and if an ISP contemplated replacing its entirety of IPv4 services with IPv6, then its IPv4-only customers

would no longer receive a useful service. Even if the ISP were to assist all of its customer base with upgrading the entirety of their equipment to support IPv6, they could not directly communicate with the rest of the Internet still using IPv4.

For this reason the transition to IPv6 was planned as a dual-stack transition, where the initial phase of the transition would see a piecemeal change of parts of the Internet from exclusively using IPv4 to a dual protocol stack mode that supports both IPv6 and IPv4. A critical part of this dual-stack transition was the guidance that wherever possible a dual-stack device would first try to establish a connection using IPv6, and use IPv4 as a fall-back. As a consequence, it was envisaged that as the amount of dual-stack capable parts of the Internet grew in size, the relative volume of use of IPv6 would rise, and the relative volume of use of IPv4 would fall. At the closing phases of the transition service providers would see so little use of IPv4 compared to IPv6 that the costs in continued support for IPv4 would outweigh the benefits associated with the residual use of the protocol, and we would see IPv4 naturally recede from the Internet. However, a second, and equally critical part of this envisaged dual-stack transition was that the transition would be essentially complete, or at worst within a very short timeframe from completion, prior to the exhaustion of IPv4 address supplies. If the transition was complete prior to IPv4 address exhaustion then exhaustion is no longer a critical issue, and if exhaustion was to occur slightly sooner, then the remaining IPv4-only parts of the Internet would be exposed to a natural incentive to convert to a dual-stack network and thereby be able to use IPv6 to communicate with the remainder of the network.

Our experience with this dual-stack transition has not followed this plan. The level of deployment of dual-stack capable networks remains very low, and we are now experiencing exhaustion of IPv4 address supplies in many parts of the world. Service providers now have to contemplate using IPv4 address compression technologies, principally in the form of NATs simply to meet the continued demands for growth in services, and IPv6 is not a viable alternative as there is just insufficient dual-stack deployment to allow an IPv6-only network service to be viable in the larger Internet. There is simply too high a proportion of the Internet that remains in an IPv4-only state at present.

This has placed an additional burden on the ISP sector that is affected by this hiatus in the supply of IPv4 addresses. These actors now face the imposition of additional capital and operational costs associated with the provision of CGN-based services in IPv4, while also still having to accommodate provision for the capital and operational costs of transitioning the network to offer dual-stack services. For many actors the first activity, that of using various forms of CGNs in their IPv4 service offerings, is a forced imperative, while the second activity has a higher degree of discretion. The case for dual-stack deployment leading to timely adoption of IPv6 is, somewhat paradoxically, also devalued to some extent, particularly if the rationale for the IPv6 deployment included the avoidance of exhaustion of IPv4 addresses. Once a provider has incurred the initial costs of the deployment of CGN technologies in its network, the marginal costs of expanding this function are, relatively, far lower.

While the level of use of IPv6 in the Internet remains low, the inherent value of the IPv6 Internet is also low. As Metcalf's law⁷ puts it, the value of a network is proportional to the square of the number of users of the network, and right now the value of the IPv4 Internet with its current user population estimated to be in excess of 2 billion⁸ is many times greater than the value of the current IPv6 network with its small number of millions of users.⁹ In Metcalf's terms the value of today's IPv4 Internet is over a billion times greater than its IPv6 counterpart

In order to see IPv6 deployed as the mainstay of tomorrow's Internet, a population of IPv6 users who number in the hundreds of millions or more is needed, which will provide a stronger business rationale for ISPs and their customers to support IPv6 as part of a dual-stack service portfolio. There are a number of potential candidates for this new wave of IPv6 users. Some commentators see the massive numbers of devices that are projected to use the 4G mobile network technology as being an ideal candidate

for IPv6 deployment, such as, for example Verizon in the United States,¹⁰ while others see the mass production consumer electronic industry and the so-called "Internet of Things" as being a major catalyst for an uptake in IPv6 in the Internet.

This report is an overview of the Internet in transition, and does not explore in detail the particular issues within selected activity sectors, such as mobile services and personal devices, wire-line access services, and content services.

Address management and IPv4 runout

This section describes the framework used to manage the distribution and registration of addresses on the Internet, and provides an update on the status of address pools managed under this framework. A further analysis of IPv6 statistics can be found in the 2011 and 2013 editions of the OECD Communications Outlook, and previous reports (OECD, 2009), *Internet Addressing - Measuring Deployment of IPv6*.

Addresses and address structures

Addresses in this context are defined by a communications protocol. The address is a unique identifier used by the communications protocol to distinguish active elements, so that the protocol can provide the necessary support to enable the communication of data between endpoints of the network defined by the operation of the protocol.

In the context of the Internet protocol an "IP address" is a unique identifier assigned to a computer's interface that runs the internet protocol suite (IP) and is connected to an IP network. The common format of an IP address is a 32-bit number, allowing for a theoretical maximum of some 4.3 billion values for unique addresses.

The IP address has a minimal internal structure, which is a division into a network part and a host part. All hosts connected to a common network share the same value of the network part of their address, and uniquely identify themselves by the unique host part of the address.

Very early implementations of IP defined an 8 bit network part and a 24 bit host part, which was then further refined into a three level hierarchy of a class (A, B, C) and within each class there was the network and host division. In Class A addresses the network part was 8-bits in length, Class B used 16-bits for the network identifier and Class C used a 24-bit network identifier part. This further division also proved to be excessively inefficient and in 1993 the address plan reverted to a simple two level network and host part, but now the network part could be any length, but had to be explicitly specified⁹. This is the current format of the addresses used by version 4 of IP. As a host address, an IP version 4 (IPv4) address is a unique 32 bit integer value. As a network identifier, the network address is also a 32-bit value, and a length value. The length value is the number of bits that are used to delineate the network part in the address.

Subsequently, version 6 of the IP protocol was specified. The internal structure of an IPv6 address is subtly different to that of IPv4. In this case the address uses a fixed length (64-bit) host identity part that is used to identify the computer's interface that is running the IPv6 protocol. But this fixed length host part of the IPv6 address does not imply that the network part is also 64 bits. IPv6 introduces a subnet field, which has a variable length of up to 64 bits. The remainder of the address part is the network identifier. As with IPv4, IPv6 networks are specified as a 128 bit value, with a length value of between one and 64. The length value is the number of bits that are used to delineate the network part in the address. The remainder of the bits up to bit 64 is the site's subnet identifier value, and the remaining 64 bits are the host identifier.¹¹

Within the IPv6 address structure there are 2^{128} unique address values (340,282,366,920,938,463,463,374,607,431,768,211,456 such values in decimal notation), but only 2^{64} (18,446,744,073,709,551,616) unique subnet values. With a common site identifier value of 48 bits, this implies that the IPv6 address plan encompasses a theoretical maximum of 282 trillion end sites (281,474,976,710,656 in decimal notation). While this is still a very large number, it is of course far smaller than the 2^{128} value range of the underlying 128 bit address space.

Description of the address distribution framework

The Internet Protocol is an open standard developed by the Internet Engineering Task Force (IETF), and the protocol parameter registration functions for the IETF are performed by the Internet Assigned Numbers Authority (IANA), and activity performed by the IANA functions operator, which is currently an activity performed by the Internet Corporation for Assigned Names and Numbers (ICANN) under terms of a contract with the United States Department of Commerce.¹²

The IANA functions operator manages the central pools of IPv4 and IPv6 addresses. These central pools are used as the distribution point for the Regional Internet Registries. In addition, the IANA functions operator manages a special purpose address registry, according to policies developed and documented by the IETF, reserving and allocating address space for special purposes related to supporting experiments and specialised address assignments relating to protocol standardisation.

For conventional use of IP addresses, the further distribution of addresses is managed by the Regional Internet Registries. There are five such registries, covering the United States, Canada and many Caribbean and North Atlantic islands (ARIN), Europe and the Middle East (RIPE NCC), the Asia Pacific region (APNIC), Latin America and the Caribbean (LACNIC) and Africa (AFRINIC).¹³ These regional registries further distribute addresses to various local Internet registries and Internet Service providers, who then use these address blocks in their networks and for customer assignment.

Statistics on address distribution

The IANA functions operator handed out its last IPv4 address block in February 2011. Some 592,708,864 addresses remain registered with the IANA functions operator as "IETF Reserved" address blocks. A further 20,466,432 addresses are being held by the IANA functions operator, representing returns to the registries of address blocks that were originally allocated to address holders prior to the introduction of the RIR system. These addresses will be evenly distributed back to the RIRs in the future in accordance with a global policy adopted in all the regions of the RIR system and ratified by ICANN as the IANA functions operator.¹⁴

The remainder of the IPv4 address space is held in the RIR registries. As of the 8th November the distribution of IPv4 addresses by registry is as follows:

Table 1. IPv4 address holdings by regional internet registry (August 2013)

RIR	Total	Assigned	Available	Reserved
AFRINIC	117,154,816	53,322,112	61,071,104	761,600
APNIC	868,197,632	850,199,992	14,100,480	3,996,160
ARIN	1,729,301,248	1,693,108,736	31,267,072	4,925,440
LACNIC	186,687,744	150,122,240	36,132,352	433,152
RIPE NCC	780,443,648	764,273,784	14,647,040	1,522,824
IANA	613,182,208	6,912	20,466,432	592,708,864

Source: Regional Internet Registries Extended Statistics Reports (www.nro.net/pub/stats/nro/delegated-extended)

In the case of APNIC and the RIPE NCC, the level of available IPv4 address space is now below that of a /8 (16,777,216 addresses). According to the address distribution policies adopted in both these regions, the address allocation policy now permits a maximum allocation of 1,024 addresses to any single applicant.

For IPv6, some 87.5% of the IPv6 address space remains reserved by the IETF. Of this total reserved space some 0.9% has been reserved for local use contexts (so called "ULA" prefixes) and 0.4% for multicast (group addresses). The remaining 12.5% of the IPv6 address space has been allocated for use in the global unicast IPv6 Internet.

The IPv6 address space held in the RIR registries is shown in the following table (the unit here is in units of /48 subnet prefixes):

Table 2. IPv6 Address holdings by regional internet registry (August 2013)

RIR	Total (/48s)	Assigned (/48s)	Available (/48s)	Reserved (/48s)
AFRINIC	68,753,031,168	297,664,978	68,057,298,880	398,067,310
APNIC	69,927,501,823	2,971,936,079	58,713,725,340	8,241,840,404
ARIN	68,887,248,896	1,834,565,020	32,512,825,981	34,539,857,895
LACNIC	68,753,031,168	417,023,843	67,913,101,621	422,905,704
RIPE NCC	70,934,003,712	3,083,274,166	59,099,570,656	8,751,158,890

Source: Regional Internet Registries Extended Statistics Reports (www.nro.net/pub/stats/nro/delegated-extended)

"Used" and "unused" addresses

When addresses are "assigned" to a local registry or ISP, the addresses can be used in a number of ways. For an ISP, or a data centre operator, it may well be the case that the address is directly used to support interactions over the public Internet. If this is the case then it is necessary that reachability to this address be advertised on the Internet's routing system, so the network will be announced into the Internet's routing system. Each individual routing announcement covers a span of addresses, so it is not necessarily the case that each and every individual address covered in the network span is uniquely assigned to an individual end device, but overall the announcement of a network address prefix in the Internet's routing system is commonly interpreted as a strong indicator that the addresses in that span are being "used" in some fashion.

The corollary is that if a network block is not announced to the global routing system, then this may be interpreted as an indication that the address block is no longer in use. This is not necessarily the case, as addresses are assigned for various purposes, including, but not necessarily limited to, their use in the public Internet. So the lack of an announcement in the Internet's routing system is not necessarily a clear indication that the address is no longer being used in an active network.

The address policy discussions within each of the regions has, from time to time, discussed the possibility of attempting to reclaim and reuse those addresses that are not in use. This option presents some issues for the registries, in so far as the agreements between the registries and address holders do not necessarily commit the address holders to immediately advertise those addresses on the public Internet, and various forms of use of addresses in contexts where the address is not advertised on the public Internet are considered acceptable forms of use. It is also the case that a significant volume of addresses were distributed by the predecessors of the RIR system, and the arrangements in place at that time placed little in the way of formal obligation on the address holder. The ability of an RIR to effectively resume such addresses without any form of active consent on the part of the address holder is generally considered to be

beyond the acknowledged powers of the registry. Address resumption efforts have had more success in those contexts where there is an active agreement between the address holder and the registry, and where the address holder has not fulfilled their obligations with respect to this agreement. However the total volume of addresses that the registries have been able to reclaim from such lapses of these address holding agreements have not been all that significant, and have had no significant impact on the address runout experience in the case of APNIC and the RIPE NCC, or in the runout projections relating to ARIN, AFRINIC and LACNIC.

The breakdown of advertised and unadvertised IPv4 addresses by each of the RIRs is shown in Table 3.

Table 3. IPv4 Advertised and unadvertised address holdings by regional Internet registry (August 2013)

RIR	Total	Assigned	Advertised	Unadvertised
AFRINIC	117,154,816	53,322,112	46,526,720	8,800,512
APNIC	868,197,632	850,199,992	726,362,300	123,765,828
ARIN	1,729,301,248	1,693,108,736	1,070,920,960	622,330,112
LACNIC	186,687,744	150,122,240	135,359,272	14,778,584
RIPE NCC	780,443,648	764,273,784	657,172,772	107,134,292

Source: Regional Internet Registries Extended Statistics Reports (www.nro.net/pub/stats/nro/delegated-extended)

Predictions on IPv4 runout

The central pool of IPv4 addresses managed by the IANA functions operator distributed its final set of address allocations to the RIRs early February 2011. At this stage it continues to manage a set of address reservations made by the IETF as part of its protocol parameter registry services function, and also manages a temporary pool of returned so-called legacy addresses prior to their re-distribution to the RIRs.

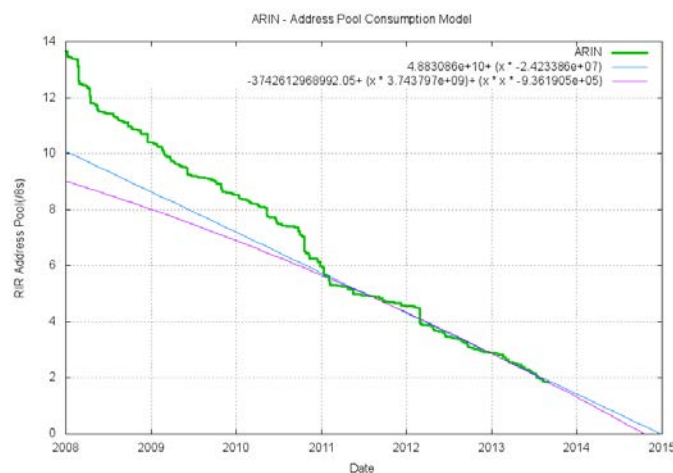
Of the five RIRs, the first to reach a conclusion of the general address allocation function for IPv4 addresses was APNIC, on 19 April 2011. Since then, APNIC is operating its continuing IPv4 allocations under a "Last /8 Policy" where each serviced entity may apply for one, and only one, allocation of up to 1024 addresses. The intent of this policy is to hold onto a small pool of addresses to assist new entrants in the area of Internet Service Provision to operate in dual-stack mode with some small amount of IPv4 to service the IPv4 side of their dual-stack needs, with the explicit awareness, noted when the Asia Pacific regional address policy community was contemplating the adoption of this particular address policy, that the address block available to each applicant under this policy could be used in conjunction with IPv4 CGNs so as to allow this very small block of IPv4 addresses to be used in far larger dual-stack networked environments.¹⁵

The second RIR to also run to the end of its general address allocation policies has been the RIPE NCC, which exhausted its pool on 14 September 2012. The RIPE NCC has also moved into a framework of a Last /8 Policy with similar constraints in place in APNIC.¹⁶

The remaining three RIRs, namely ARIN, LACNIC and AFRINIC, are yet to run out of IPv4 addresses.

In the case of the ARIN registry, it currently holds some 31,267,072 addresses in its local address pool. In estimating the projected run-out time for this registry it is noted that the model of address consumption in the region served by this registry changed significantly at the same time as the IANA registry was depleted in February 2011, which was also the time when this registry changed from a policy framework that assigned addresses to entities in a quantity that encompassed their planned requirements for the forthcoming 12 months to one that encompassed only three months of future requirement.

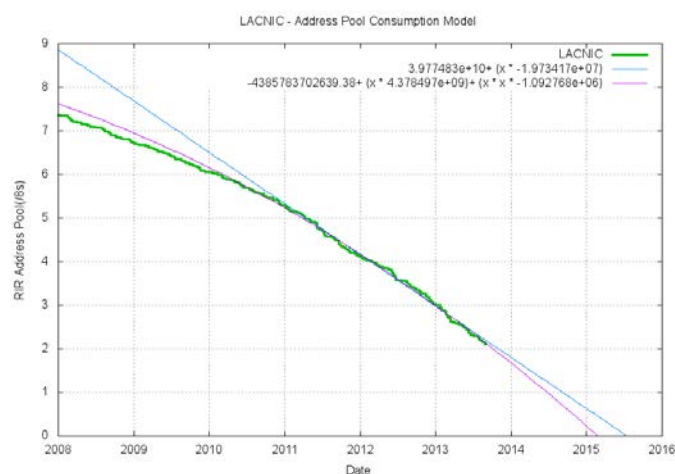
Figure 1. ARIN IPv4 address pool size, and models of address consumption (August 2013)



Source: www.potaroo.net/tools/ipv4/arinn-pool.png

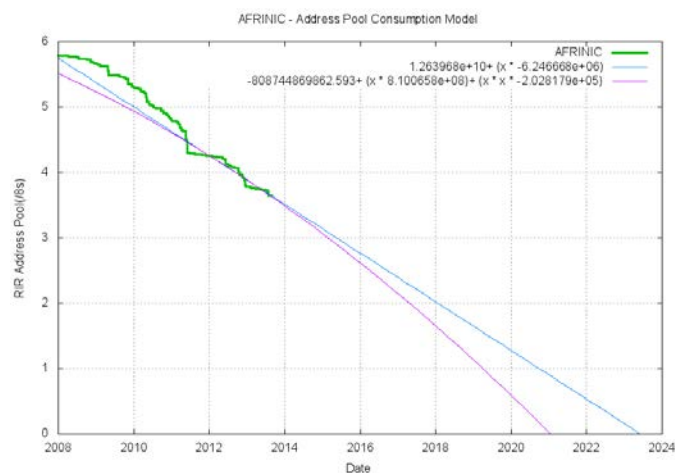
As shown in Figure 1, the demand rate for IPv4 addresses in the ARIN region is relatively constant. A study of the address consumption rates in ARIN region, undertaken in August 2013, predicts that with a 90% level of confidence, ARIN will exhaust its remaining IPv4 addresses between May 2014 and May 2015. At a 50% confidence level, the exhaustion interval for ARIN is from 19 August 2014 to 15 January 2015, with a median date of the 19th November 2014.¹⁷

The LACNIC registry is currently holding some 36,132,352 IPv4 addresses, The address allocation rate is some 20 million addresses per year, and the current estimate is that if this rate continues then LACNIC will exhaust its pool of addresses in mid-2015, using a linear estimate of future address consumption (Figure 2).

Figure 2: LACNIC IPv4 address Pool size, and models of address consumption (August 2013)

Source: www.potaroo.net/tools/ipv4/lacnic-pool.png

The AFRINIC registry is holding some 61,071,104 IPv4 addresses, and the current average IPv4 address assignment rate is some 6 million addresses per year. At this rate the AFRINIC address pool will be fully depleted in 2023, a decade from now. The rate of address consumption in Africa has been slowing down since 2009, and the predictive model is again based on a linear projection of address consumption (Figure 3).

Figure 3: AFRINIC IPv4 address pool size, and models of address consumption (August 2013)

Source: www.potaroo.net/tools/ipv4/afrinic-pool.png


There are a large set of assumptions behind these predictions, most notably that the factors that have been driving address consumption in the past will continue to drive address consumption in the future. The predictions can only act as a very approximate guide to when exhaustion of the address pool may occur in each of the three regions, but they serve at the least to highlight that the pressures brought on by address exhaustion will be experienced at different times in each of the regions.

Network address translation as a response to IPv4 address exhaustion

One response to the prospect of IPv4 address exhaustion has already been in place for more than a decade. The model of assigning a single IPv4 address to a customer and having the customer premises equipment then use a network address translating (NAT) function to share that single address across the IP devices that reside in the customer's local network is now a very well established model. In this case all the customer's devices, which are uniquely addressed within the customer's local network share a single external IP address when they communicate with external services. This model is now so well established that there is a specialised protocol for an application to negotiate a particular desired form of NAT binding,¹⁸ and many applications are well versed in detecting the presence of a local NAT and adapting their behaviour to match the variable characteristics of the NAT that is being used.

While NATs at the edge of the network are now commonplace, the key change in response to the current issues relating to IPv4 address depletion is deploy NAT functionality inside the service providers' network. This configuration potentially places two NATs in the path of a packet: the first will share a single provider's (private) address across the devices connected to the customer's network, while the second device will share a pool of public addresses across a larger pool of privately addressed customers within the service provider's network (Figure 4) (Huston, 2012a).¹⁹

Figure 4. CGN NAT model



Source: Transitioning Protocols, www.potaroo.net/ispcol/2011-03/transtools-part2.html

NATs were originally developed at the same time as the activity associated with the development of a successor protocol to IPv4. NAT was not intended as a candidate successor protocol, but was intended to be a temporary measure intended to extend the useful lifetime of the remaining IPv4 address pool. This would allow a little more time to complete the design and specification work for a successor protocol to IPv4. As a temporary measure NATs have proved to be highly successful in the Internet and are now very widely, if not ubiquitously, deployed. As a result, all applications have had to factor in the potential presence of NATs in the end-to-end path, and those that do not have had little success due to their highly limited applicability. The objective with NAT was to define a gateway mechanism that allowed globally unique IP addresses to be shared across numerous local devices. In addition, it was intended that NATs could be deployed in a piecemeal fashion within the Internet, without causing changes to hosts or other routers that did not have NAT-capable gateways. Other forms of address-sharing technologies, such as Dynamic Host Configuration Protocol (DHCP) address leasing, relied on intermittent connectivity where the address was given back to the pool when the connection was broken, whereas NATs were intended to

allow a collection of always-connected devices to share an address pool dynamically. The original specification portrays this approach as being a measure that could provide temporary relief to IPv4 address shortage while other more complex and far-reaching solutions are developed.

NATs are active units placed in the data path, usually as a functional component of a border router or site gateway. NATs intercept all IP packets, and may forward the packet onward with or without alteration to the contents of the packet, or may discard the packet. The essential difference here from a conventional router or a firewall is the discretionary ability of the NAT to alter the IP packet header before forwarding it on. NATs are similar to firewalls, and different from routers, in that they are topologically sensitive. They have an "inside" and an "outside" and undertake different operations on intercepted packets depending on whether the packet is going from inside to outside, or in opposite directions.

The advantages of NAT are that NATs do not require any changes to end hosts or routers. Whether a NAT is in place between the local network and the Internet or not, local devices can use the same software and support the same applications. As long as one can accept the limitation that sessions must be initiated from the "inside," NATs can work in an entirely transparent fashion for a large set of client-server classes of applications. NAT represent an effective, provider-independent addressing solution with multi-homing capabilities. NAT allows for rapid switching to a different upstream provider, by renumbering the NAT address pool to the new provider's address space. In essence, NATs provide the local network manager with the flexibility of using provider-independent space without having to meet certain size and use requirements that would normally be required for an allocation of public, provider-independent address space. Most critically, NATs do not require a co-ordinated deployment. There is no transition, and no "flag day" across the Internet. Each local network manager can make an independent decision whether or not to use a NAT. This allows for incremental deployment without mutual dependencies.

However, NATs represent a set of design compromises, and no examination of NATs in today's Internet would be complete without noting some of their shortcomings. NATs cannot support applications when the initiator of a transaction lies on the "outside". This implies that peer-to-peer services, including conventional services that emulate telephone services, or novel services, such as distributed data caches and highly replicated data distribution services, cannot function in a NAT environment without the use of middleware applications (or "gateways") to assist the application. This workaround introduces a modification to the service architecture, where forms of "any-to-any" services require the deployment of a distributed set of deployed "agents" and "helpers" to dynamically translate the application level identifiers into transport IP addresses. This adds cost and complexity into applications, and adds interdependencies across discrete network elements. The behaviour of NATs varies dramatically from one implementation to another. Consequently, it is very difficult for applications to predict or expose the precise behaviour of one or more NATs that may exist on the application data path. In summary, workarounds to support services that function robustly across NATs are often limited, complex, and fragile.

NATs also present issues at the network level. NATs may drop IP packet fragments in either direction: without complete TCP/UDP headers, the NAT may not have sufficient stored state to undertake the correct header translation. NATs have no inherent failover. NATs are an active in-band mechanism that cannot fail into a safe operating fall-back mode. When a NAT goes offline, all traffic through the NAT stops. NATs create a single point where fates are shared in the NAT device maintaining connection state and dynamic mapping information. NATs sit on the data path and attempt to process every packet. Therefore bandwidth scaling within the network requires NAT scaling at its boundaries.

A further consideration relates to the level of address "compression" that is used by the NAT. The basic NAT function when a single IP address is being shared across multiple internal devices is to use the same source address for all outbound packets, but use a unique TCP or UDP source port address for

each unique transport session that is concurrently supported on the NAT. This typically allows a single IP address to be shared across some hundreds or at most some thousands of customers when used in the context of a CGN. However a far higher address compression factor can be achieved by allowing the same externally visible mapped TCP or UDP source port address to be used concurrently by different external transport sessions. This extends the theoretical maximum number of concurrent NAT sessions from 131,072 sessions (16 bits of port addresses in each of TCP and UDP transport protocols) to a theoretical maximum of 2^{64} sessions. While this theoretical maximum is completely unachievable, this NAT technique effectively allows a single IPv4 external address to be shared across thousands to as high as tens of thousands of internal users.

The introduction of NATs into the carrier's architecture also raises the issue of the cost and complexity of data logging. Many regimes apply some form of data retention directive to ISPs. Typically this directive calls for the carrier to retain their DHCP logs, which effectively records the identity of the customer identity and the IP address assigned to the customer, and the time that the assignment was made and the time the IP address was handed back into the DHCP pool. The data requirement to support this directive is of the order of 15Gb per month per million users, which is not a highly onerous imposition even with data retention periods of 2-3 years. This is not necessarily the case for CGNs, where the logging requirement now requires not only the mapping of internal to external address and port, but also the destination address and port, the transport protocol identity, and an accurate high definition time of the session start and the time of the session completion. A study by CableLabs²⁰ of the data retention requirements for a CGN points out that a comparable data retention requirement in a CGN needs to log a connection record for every single transport session made by a user. CableLabs reported that the current connection volume is some 33K to 216K connections per customer end point per day, and allowing some 450 bytes per connection record and an 8 hour busy period, the same million customers would generate a CGN log stream that would peak at 30 Gbps, and 2.9 Pb of storage requirement per month. This represents a considerable cost and a considerable challenge to any form of search across the retained data using existing technology capability

The costs involved in the use of CGNs have been recently modelled.²¹ This study estimates the cost of a CGN to be of the order of USD 90 000 in capital expenditure to serve 10 000 end users, and an additional USD 10 000 p.a. in operational support costs. This study also included an estimate of customer support costs, and additional customer churn to estimate a total cost of the use of CGN to some USD 40 per customer per year. The major element of cost here is the consideration of customer churn and lost revenue arising from this churn. If all service providers were using CGNs the churn factor would be mitigated significantly, and the resultant cost of CGN ownership is estimated to be some USD 3.00 per customer per year based on the data presented in this study. This study uses a theoretic approach and is not based on actual experience, so the estimates of the support costs in offering a dual-stack service represent the greatest level of uncertainty in this study. It is also noted that this study did not factor in any data retention costs. With a higher level of support required, together with an estimate of data retention costs, it is possible to estimate a per-user annual cost of USD 60 p.a. without factoring in the customer churn consideration. It is not easy to foresee whether these costs will rise over time. Higher levels of NAT compression would see NATs with larger memory tables and faster processor speeds in order to service the same throughput rates, but at the same time the downward pressures on price due to continual refinements in technology speed and capacity may well offset such factors.

NATs were a short term expedient measure that has turned into a longer-term set of overriding constraints imposed on the further evolution of the IPv4 Internet. Not only do novel applications in today's IPv4 Internet need to include considerations of NAT traversal, but we appear to be entering into a situation where if an application cannot work across NATs, then the application itself fails to gain acceptance. We seem to be locking into a networked world that is almost the antithesis of the Internet concept. In this NAT-based world, servers reside within the network and are operated as part of the service provider's role,

whereas end devices are seen as "dumb" clients, who can establish connections to servers but cannot establish connections between each other. The widespread use of NATs appears to be reinforcing a re-emergence of the model of "smart network, dumb clients," whereas others would argue that the network is getting no smarter, it is just that the number of obstacles and amount of network debris is increasing while clients are getting worse at maintaining coherent end-to-end state in the face of such changes.

The IPv6 protocol

The Internet protocol suite is an instance of a layered protocol design. At the lower levels the IP protocol interfaces to certain link level media types, while the IP layer interfaces to the upper layer transport protocols. The IPv6 protocol does not change the function or operation of the Internet protocol suite end-to-end transport layers. If an application that uses an interface to the network at a socket-level abstraction to the network would be expected to operate equally well over an IPv6 network as an IPv4 network. The implication is that the world of applications on the Internet, at least from the standpoint of the applications' interface to the network, do not notice any difference between these two protocols.

The changes between the IPv4 and IPv6 protocols are changes specifically relating to the datagram functionality of the network, and the associated address framework. There are no changes in the end-to-end transport protocols of TCP and UDP. Here there are some network level changes, some of which are substantive, while others are of less significance.

(a) What changed with IPv6

Obviously the address fields in the IP packet header have changed with IPv6. These fields have expanded from 32 bits to 128 bits in length.

The other significant change in this transition has been the management of fragmentation in IPv6. IPv6 routers are unable to fragment a packet that is too large to be forwarded onward. In IPv6 the router must generate a IPv6 Internet Control Message Protocol (ICMP6) message back to the sender to indicate the problem, and to inform the sender what packet size would be acceptable. Within a reliable end-to-end transport protocol, such as TCP, this return-to-sender function, while a change to the fragment-on-the-fly behaviour of IPv4 is not a major shift in protocol behaviour. However, if the upper level application uses UDP, then the ICMP6 message does not have the capability to generate a retransmission of the offending packet, in which case the repair has to be made by the upper level application.

Other changes relating to the handling of the IP options field and the introduction of the IP flow label have no significant impact on the semantics of the IP protocol.

(b) And what did not

There are no other substantial changes in the IP protocol with this change from IPv4 to IPv6. The basic architecture of IP remains constant, with the IP model defining a simple datagram delivery service layered on top of a packet-based media adaptation layer. The end-to-end transport layers remain unchanged.

The routing protocols have remained largely unchanged, although some routing protocols have required minor modification to carry IPv6 prefix reachability in the payload of the protocol.

At the same time, the security capabilities of the protocol remain largely unchanged.

From this respect, IPv6 represents a conservative incremental change to the architecture of the Internet Protocol. It has basically changed the length of the address fields in the IP header, and also changed the packet fragmentation behaviour, but has changed little beyond these two characteristics.

(c) The demand drivers for the IPv6 transition

Much has been said about the search for a highly desirable application that would require IPv6 in order to function. This has been termed “the search for the killer app,” and has been a constant theme of many IPv6 conversations. To date there has only been one consistent response to this question: the major driver for IPv6 lies in its larger address space, and its ability to avoid the need of middleware to allow end devices to connect directly to the Internet.

The changes to the packet header structure in IPv6 do not have any major impact on the cost of switching IPv6 packets. Within the network's packet switching infrastructure the functional requirements to extract the address, hop count and service class fields from an IP packet and use that to perform a switching and queuing function are little different between the two protocols, and, compared to the economies of scale of production, are not a differentiator for a service provider.

The elimination of the absolute need for middleware is not in and of itself a major demand driver for IPv6. Today's end users often assume the extensive use of filtering functions performed by middleware as part of the security framework for Internet services, and as such it would be unrealistic to anticipate that an all-IPv6 network would eliminate the continued demand for middleware to be deployed as part of the network's infrastructure.

It would also be unrealistic for application designers to assume that all communications are open clear channel transactions. Application designers make use of various techniques to sense the environment in which they have been invoked, and, as required, adopt a mode of functioning that is intended to be effective in that environment. For example, applications that are sensitive to the presence of NATs on the path may use a NAT sensing protocol to determine the nature of the NAT and adjust their behaviour to suit the NAT function.

However, these factors are all factors in an environment that are looking for demand drivers that differentiate IPv6 in today's Internet, and are drivers that are relative factors in a network that makes use of IPv4 and various forms of NATs. The long term issue is that continued growth of the Internet will place inexorable pressure of the NAT function, and the need for ever larger factors of address “compression” may ultimately fail in the face of an ever larger population of communicating devices. The longer term demand driver is that the IPv4 Internet will be unable to sustain the communications requirements of a more populous Internet that we confidently anticipate in the coming years.

However, long term demand drivers and short term investment cycles can find themselves in a disjointed configuration. In order to stimulate short term demand in the market we have seen a number of programmes adopted in the public sector that promote the use of IPv6 by public sector agencies. Some examples of this form of stimulating the market in IPv6 services include the public sector initiatives in the United States,²² Australia,²³ Sweden,²⁴ Germany²⁵ and Japan.²⁶ It is anticipated that these initiatives will stimulate domestic demand for IPv6 services and assist in speeding up the transition efforts. While there may not necessarily be exclusive causation here, it is noted that the level of use of IPv6 in the United States, Germany and Japan have increased significantly in recent months, and the public sector programmes in these countries may be a contributory factor here.²⁷

(d) *Backwards compatibility issues*

Despite the relatively small changes in terms of the difference between the IPv4 and IPv6 protocols, it is nevertheless the case that an end device that uses only IPv4 to communicate cannot initiate or receive a communication from another device that uses only IPv6. In other words, IPv6 is not "backwards compatible" with IPv4. IPv4 only devices can only communicate with other devices using IPv4, and, similarly, IPv6 devices can only communicate with other devices using IPv6.

An IPv4 protocol instance will only accept packets formatted using the IPv4 specification. The IPv6 specification is not implemented as IPv4 together with additional information placed into option fields. It has been implemented using an entirely different packet header format. The consequence of this is that an IPv4 host will be unable to process any received IPv6 packets.

This has a number of implications in terms of the transition of the Internet from an all-IPv4 network, through a hybrid phase of supporting both IPv4 and IPv6 in a "dual-stack" mode, through to the ultimate objective, that of an all-IPv6 network.

The IPv6 transition plan

Given the lack of backwards compatibility, there is no simple procedure to transition to an all-IPv6 Internet via unco-ordinated piecemeal steps. It is not possible to variously upgrade the components of a network from IPv4 to IPv6 and provide a constant set of connectivity services through the upgrade process. Instead, the network has to undertake a more complex form of transition that involves a transitory intermediate phase.

The transition process being used for the deployment of IPv6 in the Internet is a so-called "dual-stack" transition. What is possible is to support the various upgrades of the components of a network from IPv4-only to a dual-stack network that supports the active use of IPv4 and IPv6 concurrently. This mode of operation is a "dual-stack" mode, where devices that support both protocol stacks have the choice of being able to use either protocol.

The way this dual-stack mode of operation is supported within an application, such as, for example, a web browser, starts with a URL or a similar name-based rendezvous point to start the network transaction. If the local host has an active IPv4 interface it will query the Domain Name System (DNS) for an IPv4 address for the DNS name component of the URL. If the local host has an active IPv6 interface it will perform a similar DNS query, this time for the IPv6 address that has been bound to the given domain name.

If the DNS queries result in only an IPv4 or an IPv6 address for the named service point, then the application will initiate the connection in the protocol that corresponds to the available address. If the DNS query returns addresses in both protocols then it will use a local preference rule set to determine which protocol to use.

Initial implementations of this local protocol rule set preferred IPv6 over IPv4 in most cases. The implication of this on the overall dual-stack transition is to create a positive feedback on the use of IPv6. In the initial phases of the transition it would be expected that few DNS queries would result in both IPv6 and IPv4 answers, so the Internet would remain a largely IPv4 communications domain. As the number of services using both IPv4 and IPv6 increased in number, then the number of DNS queries that result in both IPv6 and IPv4 increase, and by virtue of the local preference roles, the number of connections made using IPv6 would increase and the number of IPv4 connections would correspondingly decrease. Ultimately if the overwhelming majority of systems were configured as dual-stack systems then the result of this procedure would be that the majority of connections will be made using IPv6. At this point the requirement for

universal support of IPv4 would wane, and it is logical to anticipate the switching off of IPv4 at this juncture.

Given that the IPv4 and IPv6 worlds cannot be bound together in a seamless fashion at the protocol level, then the intermediate waypoint of the transition is the objective that the network needs to simultaneously support two distinct protocol families: IPv4 and IPv6. As the amount of dual-stack services increases, the level of use of IPv6 is anticipated to rise, at the expense of the level of use of IPv4.

Presumably, it would be known that the final stages of this transition had been reached when dual-stack networks ceased to use IPv4 at all, which would occur when all the remote counterparts in their communications transactions were also configured as dual-stack systems. At that stage the continued need for IPv4 would have finished.

This broad plan raises more questions than it answers, particularly relating to the nature of the drivers that would see existing deployments of IPv4 infrastructure convert to a dual-stack environment, and to the drivers that would see the culmination of the transition and termination of the dual-stack environment. There is no particular timetable associated with the plan, and no particular means of co-ordination of activity, so that knowing when we have completely finished with IPv4 may be more challenging to determine than the simple procedures described here.

(e) Assumptions

The major assumption behind the dual-stack transition approach is that both protocols are equivalent, and that the service delivery experience as perceived by the end user will be independent of the protocol used to deliver the service.

If this is indeed the case, and that the two protocols experience the same level of quality in the delivery of service, then a local configuration setting for dual-stack clients that prefers to use IPv6, namely to "try IPv6 first, then fail over to IPv4" would be a viable approach.

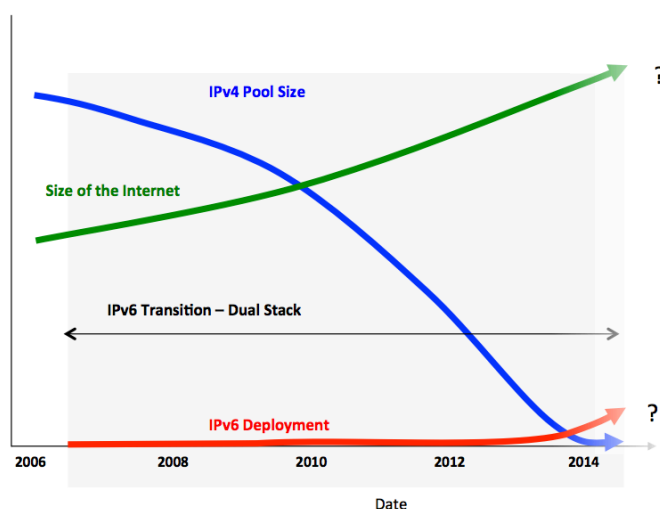
This is not, however, the entirety of our experience so far, and the general observations of service quality in the two protocols point to some lingering issues with the quality of the service provided via IPv6. Measurements of the connection failure rate in IPv4 point to a very low residual connection failure rate at a level of 0.01% or thereabouts (at this level the signal of connection failure is lost in the actions of malware SYN scanners, and similar malware behaviour), while the connection failure rate of IPv6 connections across today's network is considerably higher at some 2% of all IPv6 connection attempts. This has negative implications for the "IPv6 first" dual-stack connection approach, as the failing connection attempts take variously between 21 and 109 seconds to resolve and fall back to IPv4.²⁸

The assumptions that the quality of the IPv4 and IPv6 networks are comparable, and the assumption that failover from IPv6 to IPv4 is without visible cost or penalty are both turning out to be flawed assumptions. There is, however, a further assumption in the original dual-stack transition plan that is also turning out to be flawed, and this additional factor may well be the most critical one in the entire dual-stack transition. This assumption was that the transition would be undertaken during the period while there was still abundant IPv4 addresses, or at a very minimum, sufficient IPv4 addresses to allow this transition to run to its completion without requiring separate measures to make more efficient use of the available IPv4 addresses.

(f) *Dual-stack and IPv4 address exhaustion*

The exhaustion of the IPv4 address pools are proving to be a major complication for the transition. The original assumption was that the dual-stack transition would take place while there was still adequate supplies of IPv4 addresses to sustain the demands of a still expanding and still highly competitive environment in the supply of IP carriage and access services. With the exhaustion of IPv4 addresses much of this original plan has to be modified (Figure 5).

Figure 5. Dual-stack transition plan - with IPv4 exhaustion



Source: G. Huston. "IPv6 Transition Plan", October 2010, www.potaroo.net/presentations/2010-10-19-ipv6-transition.pdf

If a service provider is operating a service that continues to expand, then the service provider is now facing the additional business cost of having to invest in some form of CGN technology in addition to the dual-stack platform costs.

There is no single technical approach being used in this situation. While a conventional approach is to use private IPv4 address space within the service provider network and a second NAT function to map the internal private addresses into public addresses, other approaches also exist. The NAT64 approach proposes the exclusive use of IPv6 services within the service provider's network, and the use of protocol translation to map the internal IPv6 addresses to an external IPv4 address in the case that the external party is IPv4 only. More complex models perform this translation both at the interface to the public network and also to the customer's edge, so that the customer's edge network is a dual-stack IPv4 and IPv6 network, and the external capabilities are also both IPv4 and IPv6, but the internal service network is an IPv6-only network. There are a number of alternative approaches, including mechanisms that use address plus port mappings and Dual-Stack Lite²⁹ that uses IPv4-in-IPv6 encapsulation rather than having a second level of IPv4 network address translation. The variables here are whether to translate or perform encapsulation, whether to operate the service provider network in single protocol or dual-stack protocol mode, whether to use altered CPE equipment or try and leave the CPE functionality unaltered, or require changes to the CPE model.

All this adds to the level of variability in transition structures and this variability becomes an issue in terms of a standards based environment that uses tested interoperable components. The evident alignment of transition technology to individual vendors points to the emergence of a highly customised set of options for service providers where there is a set of compromises between the choices available for the network infrastructure and choices available in the edge equipment that interfaces the network to the customer's local network.

There is no single optimal strategy for the combination of dual-stack transition and IPv4 address exhaustion, and there are a number of variables here that add to the complexity and cost of this transition process.

The first variable is that there is no common view as to how long this transition will last. This means that when a service provider makes an investment in some form of CGN technology there is no clear model of when the level of address utilisation is as desired. There are a number of different NAT strategies, and they vary in terms of the level of simplicity of NAT operation and the level of address utilisation, and these variances will become increasingly critical in scenarios of extended transition periods.

The second variable is that there is no single transition strategy. This means that across the entire set of Internet service providers there will be a variety of transition models, each of which will induce various forms of network behaviours under certain conditions. For the application designer this raises the issue of how to engineer a robust application that will function reliably under all possible circumstances, where the range and behaviours in these circumstances are difficult to predict.

The third variable is that there is no coherent organisation of the transition, while the transition itself requires some aspects of organised activity in order to reduce the potential for transient failure in the IPv6 service environment.

Finally, the transition to IPv6 does not offer a service provider any significant leverage in terms of efficiency of operation.

Measuring the state of the IPv6 transition

There have been many approaches over the years relating to establishing metrics to trace the progress of IPv6 adoption in the Internet. These various measurements and their associated interpretation reflect a panoply of perspectives on IPv6 in the context of the Internet, and also reflect the fact that the Internet is not a single integrated system but a coalition of component subsystems, and the measurement of IPv6 can be performed within the context of any particular subsystem.

This section briefly covers a number of these measurements and look at their assumptions in terms of being a good indicator of overall adoption of IPv6 in the Internet today.


Measurement using the routing system

The routing system is a unique subsystem in the Internet in that it provides a perspective on the entire Internet in a single view. The routing system can be used to track the number of advertised routes that constitute the IPv4 Internet, and compare this with a comparable count of the number of routes in the IPv6 protocol.

Both IPv4 and IPv6 exhibit continued network growth. Since January 2009, the number of routes in the IPv6 network has grown from 1 800 entries to the current (November 2012) value of 11 000 entries. A comparable picture of the IPv4 network shows a similar picture of growth, with the numbers rising from some 280 000 routes in January 2009 to 430 000 routes some 46 months later. There are a number of

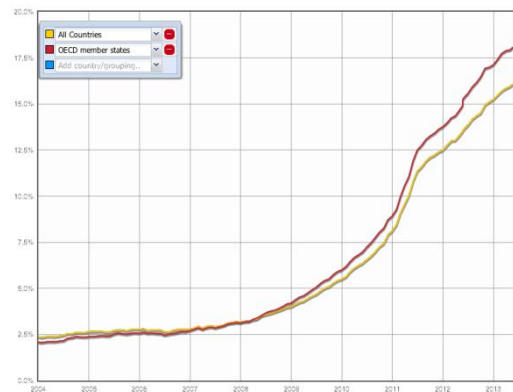
perspectives on these four numbers. In absolute terms IPv6 has risen by some 9 000 routes in 46 months, while IPv4 has risen by some 150 000 routes in the same period. The relative growth in terms of routing table entries is that IPv4 entries in the routing table have grown at a rate that is 17 times larger than the equivalent growth in IPv6. Another perspective is to look at the size of the IPv6 routing table as a proportion of the size of the IPv4 routing table. How this ratio has changed over time since January 2004 can be shown (Figure 6). Since mid-2008, the routing table size of the IPv6 network has been growing at a faster rate than that of the IPv4 routing table size, and since January 2009 the IPv6 routing table has grown in relative terms from 0.55% of the size of the IPv4 routing table to the September 2013 value of 3% of the size of the routing table.

Figure 6. Ratio of IPv6 advertised prefixes : IPv4 Advertised Prefixes



Source: <http://bgp.potaroo.net/stats/nro/v6>

Of course the number of entries in the inter-domain routing table is not the only metric of deployment size of these two IP protocols. It may be more useful to look at the number of routing entities that are routing IPv6, where each autonomous routing entity, (commonly a "routing entity" corresponds to an ISP or corporate network) is counted only once. In this case it's not the number of entries in the Internet's inter-domain routing table *per se*, but the number of unique autonomous system numbers that are contained in the routing table that indicate the number of entities that have IPv6 networks of one form or another that interconnect in the global IPv6 Internet. The reason to look at autonomous system numbers rather than route prefixes is that the IPv4 routing table has a certain amount of inherited legacy of fragmentation of network announcements that is not replicated in IPv6. By directly comparing the use of the two protocols on a network-by-network level, rather than by individual announced prefixes, then there is a better assurance that the comparison is between like artefacts in the two network protocol families (Figure 7).

Figure 7. Ratio of IPv6 AS's : IPv4 AS's

Source: http://v6asns.ripe.net/v/6?s=_ALL;s=_OECD


The number of AS's announced in IPv6 has risen substantially as compared to the number announced in the IPv4 network since the start of 2009. In January 2009 the number of AS's in the IPv6 routing table was some 3% of that found in the IPv4 routing table, and this has risen to a value of 17% by September 2013. This metric paints a somewhat different, and more positive, picture of IPv6 deployment than the comparison on the number of entries in the routing table.

There are a couple of potential issues in the routing table view of the IPv6 Internet. Firstly, a metric of capability of supporting IPv6 in routing is not the same as a metric of actual use of IPv6 in terms of services on IPv6, and IPv6 packets that are sent across the network.

(g) Measurement using the domain name system

The Internet's domain name system can also provide a perspective on the state of IPv6 deployment. For a client to initiate a connection to a server using IPv6 it is necessary for the client to learn the IPv6 address of the server. This is a function of the Internet's domain name system (DNS). Can the count of the number of domain names configured with IPv6 addresses provide an insight into the level of deployment of IPv6?

One approach is to take a list of the more popular web sites and see which of them have an IPv6 address. The most common source of such popular domain names is the Alexa list (www.alexa.com), and the measurement technique is to query this set of domain names to establish what proportion of the names have an IPv6 address. One such longitudinal study was undertaken by Lars Eggart since 2007 (Figure 8).

Figure 8: IPv6 address records in Alexa Domain Names

Source: L. Eggert: www.eggert.org/meter/ipv6

What this shows is a steady uptake in a number of countries with IPv6 addresses on their domain names, with two pronounced discontinuities, occurring in July 2011 and July 2012. Globally the proportion of domain names with IPv6 addresses now sits at some 20% of the set of the most popular domain names.


A related domain name-derived metric is the proportion of clients who are capable of resolving domain names using the DNS protocol over an IPv6 transport. This is not a direct client capability measurement, but it is a measurement that reflects the degree to which the common infrastructure of the Internet, and in particular the DNS name resolution infrastructure, is capable of operating in a dual-stack mode and is equally capable of operating over IPv6 as it is over IPv4.

In September 2012, an experiment that tested the capabilities over a random sample of more than 2 million clients found that some 18% of these clients used DNS resolvers that were capable of supporting queries over IPv6.

Both of these metrics point to a relatively encouraging position with respect to the integration of dual-stack capability into the DNS infrastructure of the Internet.


Measurement using Internet traffic statistics

Another form of traffic measurement is to look directly at traffic volumes in IPv4 and IPv6. The issue with this form of measurement is that such data is generally considered to be proprietary data, and this is not released as public data. However, a number of exchange points do publish public data about the volumes of IPv6 traffic, and the Amsterdam Internet Exchange (AMSIX) shows this data as part of their public reports. The growth in IPv6 traffic over AMSIX over the past 12 months, showing a rise in traffic from an average of 1.7 Gbps in late 2011 to some 5.5 Gbps in August 2013, can be illustrated (Figure 9).

Figure 9. IPv6 traffic volumes at AMSIX

Source: www.ams-ix.net/technical/statistics/sflow-stats/ipv6-traffic

A similar report can be found at DE-CIX in Frankfurt, which allows the comparison of IPv6 and IPv4 traffic volumes seen at the exchange (Figure 10).

Figure 10. IPv6 traffic volumes at DE-CIX

Source: <http://www.de-cix.net/about/statistics/>

The daily IPv6 traffic profile at DE-CIX is peaking at some 10 Gbps in September 2013, while the IPv4 traffic load is peaking at 2.5Tbps. In this particular exchange point the traffic profile of IPv6 sits at some 0.2% of the traffic profile for IPv4.


Measurements of end client capabilities

None of these metrics presented so far are entirely satisfactory. They reflect the state of IPv6 use within the component subsystems of the Internet, but do not necessarily reflect the level to which each of these component systems interoperate and work together to deliver the necessary end-to-end service in IPv6. In looking at a typical Internet transaction, it typically starts with a URL in the form of a domain name and a service point. The first task is that of the DNS, which will be used to perform a query for an IPv6 address for the domain name. Assuming that the name has been provisioned in the DNS with an IPv6 address, then the query will return this address and the connection will move on to the next step. At this point the system will attempt to open up a connection with the remote server by sending it a packet and will await a response. For this to happen the local interface must be IPv6 capable, the local network must

be IPv6 capable, the local ISP must support IPv6 and the routing system must be carrying IPv6 routes. In other words, for a client end system to be able to make a connection using IPv6 then all of the Internet's subsystems must also be functional in supporting IPv6.

One simple way to measure the number of IPv6-capable clients in today's Internet is to use a dual-stack service point and count the number of clients who successfully establish contact with the service point using IPv6. This form of measurement has been undertaken on service infrastructure operated by Google (Figure 11):

Figure 11: Percent of end hosts preferring to use IPv6 when accessing Google's dual-stack servers



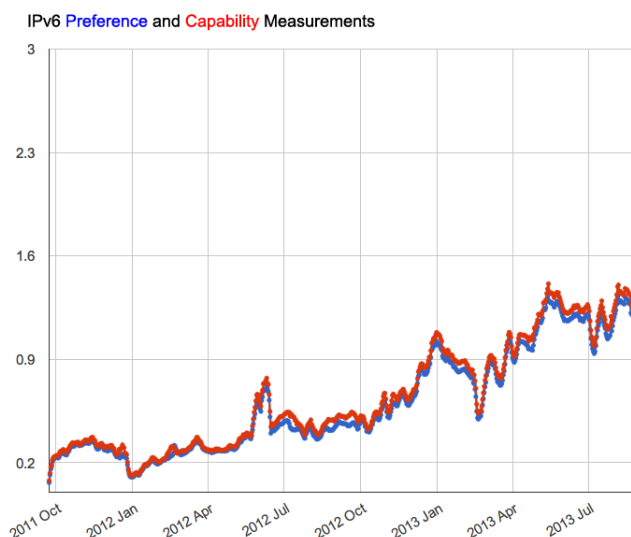
Source: www.google.com/ipv6/statistics.html

This plot shows that relative proportion of clients who prefer to use IPv6 in a dual-stack situation has doubled over 2012, rising from 0.4% of clients at the start of the year to some 0.9% by November 2012. In 2013 this has further increased, to a measurement of some 1.6% of clients by the end of August 2013.

It is possible, however, to perform a more detailed measurement of client capabilities by expanding this protocol choice exercise into a simple IPv6 connectivity "test". One such technique is to embed a relatively simple IPv6 connectivity test into a web-based resource. Such a script can pose a small number of download "tests" for the user that requires the user to load a web object that can only be loaded using IPv6, and then direct the user's browser to load a second object that is accessible using either IPv4 or IPv6 (dual-stack). If the end user can only use IPv4 then the user will fail the first test, and use IPv4 in the dual-stack test. An end user who has a dual-stack service will pass the first test using IPv6, and typically use IPv6 in the dual-stack test. "Typically" because until recently a dual-stack system would always try to use IPv6 first and then fail back to IPv4, whereas recent changes in a number of systems now use a local selection algorithm that will choose between IPv4 and IPv6 depending on what appears to produce the fastest outcome at the time. It is the effects of this recent change in dual-stack client behaviour that is the difference between the initial simple test of dual-stack preference, which assumes that in dual-stack mode clients will always prefer to try to connect using IPv6 first, and this second form of scripted test that explicitly tests the client using an IPv6-only test in addition to the dual-stack preference test. The result of this form of capability measurement when applied to some 30 million end clients per month (Figure 12).

This indicates that across the entire Internet in August 2012 an average of some 15 in 1000 end clients are capable of performing an end-to-end network transaction using IPv6 (and do not use auto-tunnels to perform this form of access).

Figure 12. Percent of end hosts capable of undertaking IPv6 network transactions



Source: <http://labs.apnic.net/ipv6-measurement/Regions/001%20World/>

Potential scenarios

At this stage it is not clear how the transitional situation to IPv6 will be resolved in the longer term.

One potential scenario is that the number of services and clients using IPv6 in this dual-stack context achieves a sufficiently large population that it becomes viable for new entrants to exclusively use IPv6 as their connection protocol without any loss of perceived utility in terms of using the Internet. A natural consequence of this outcome is that the added cost of continuing to support a dual-stack environment would have marginal utility, and we would see a decline in the level of support for IPv4, a drop in demand for IPv4 addresses and the assumption of a new role for IPv4 as a legacy technology within specialised private environments where other factors may be present.

A less desirable potential scenario is that the number of services and clients using IPv6 in this dual-stack context remains marginal, and IPv4 services remain essential to realise the utility of a comprehensive Internet service. It is likely within this scenario that the relative level of use of IPv6 will plateau and then decline, leading to an Internet that makes extensive use of various forms of middleware, and domain-based content distribution systems, and one where end-to-end services are no longer possible. The level of transformations placed on communications models in order to provide ever greater level of address utilisation efficiency imply that the Internet will only support a small number of session models that cluster around the functionality provided by web services.

Some of the implications of these scenarios, in terms of their consequences for the future of the Internet itself, can be examined.

(h) Openness and innovation

Very little of the Internet of 10 years ago is evident in the Internet of today. The last decade has seen the transformation of the mobile device from a telephone to a digital communications system. This has transformed much of the application space of the Internet, introducing services that feature immediacy, rich media, and large social networks.

One of the significant aspects of this change is that none of these innovations in the form of products, applications and services required the permission of network service providers to be launched. This has led to a decoupled collection of markets, where the market for Internet access services at a carriage level is entirely separate from the market for the provision of content and services. This second market has been commonly termed as an "over the top" market model to graphically illustrate the level of disconnection between carriage services of IP datagrams and the content services relating to the provision of various goods and services. This second market is an intensely competitive environment that focuses on the needs of consumers.

Much of the construction of online goods and services is built upon the foundations of open technologies and open networking standards. This allows applications to be constructed using existing open libraries and toolkits, allowing for rapid and highly cost efficient development of services.

(i) Carrier grade NATs

A scenario of increasingly intensive use of NATs in the Internet has a negative influence on openness and innovation. An application cannot rely on a consistent NAT behaviour, and therefore has to incur additional overheads for even basic network transactions. There are also implications in terms of server robustness under scaling pressure. A major issue here is that NATs are not constructed to any common technical standard, and do not operate in consistent and uniform ways.

In a scenario of increasing scarcity pressure on IPv4 addresses, a CGN operator may be under considerable pressure to increase the address sharing ratios of their equipment, which, in turn impacts the availability of protocol port addresses per customer, which, in turn impacts on applications. In such a scenario applications that make strong use of parallelism, such as, for example, Google's Map application, would be negatively impacted. This can force applications back into a serial model of operation, which, in turn, has a negative impact on the efficiency and speed of the application.

As a consequence, the scenario of extensive use of CGNs by network operators would reduce the flexibility of the network and increase the impediments to supporting communications that involve multiple parties, various forms of peer-to-peer connections, or applications that operate in an "always on" mode. These communications would no longer be viable as an edge based distributed application, and would require the assistance of servers and gateways within the network infrastructure. The technical implications of using NATs impact the design and flexibility of application models, and increase fragility of Internet-based services.

A number of public policy regimes support the concept of "network neutrality" where network operators are required to operate in a manner that is neutral or impartial to various application level transactions involving services and content delivery. The question posed by the introduction of CGNs into this environment is whether this neutrality can be maintained.

The second issue is that CGNs provide carriage providers with a direct insight into the upper level service transactions being undertaken by users, as every TCP session, and every UDP data stream, will trigger a NAT binding in the CGN, which will be a transaction that is visible to the carriage service operator. This log of the actions of the CGN in creating and maintaining translating table binding entries in

response to every TCP and UDP transactions becomes, in effect, a complete log of the online activities of each and every customer sitting behind the CGN. As a number of content industries have illustrated already, knowledge of consumers' actions, and inference of their preferences and habit, is knowledge that has some considerable value for advertising placement brokers and similar enterprises.

To the extent that this form of network middleware has the potential to alter the current balance between the carriage and content industries, and allow carriage providers to use the CGN behaviours to support the imposition of constraints and controls over the communications between consumers and content providers, the deployment of IPv4 CGNs in the Internet should be a matter of public policy concern.

Conclusion

The Internet has often been portrayed as an outcome of a liberalised telecommunications environment that was changing from being generally characterised as a command-led economy to one that is largely a market-driven economy. That portrayal does not mean that the interplay of market forces and the desires and motivations of market actors will always generate outcomes that strike appropriate balances between short term expediency and longer term common interest, between private interest and the public good, and between servicing the needs of incumbents while still ensuring the ability for new market entrants to innovate and challenge these incumbents.

To date, the Internet has followed a path that appears to balance these various pressures within acceptable levels of tolerance. However, it has done so in an environment of prodigious abundance. On the computing side Moore's Law continues to deliver ever more powerful processors that operate at faster speeds, while consuming less power and generating less heat. This has allowed computing processing capability to become an abundant commodity that exists in widely dispersed embedded devices as well as being concentrated in warehouses of data processing. At the same time the shift from electrical communications systems to optical systems, and the use of digital signal processors and advanced optical wave division multiplexing, coupled with polarisation phase modulation capabilities continues to transform communications capabilities, with the bandwidth of communications systems continuing to rise and the unit cost continuing to drop. This has generated a market of abundance where computing and communications capabilities are now being embedded into consumer and utility devices in a manner that is completely foreign to the dedicated computing engines of two decades ago.

Continual evolution and innovation is fuelled by the abundant capacity of essential inputs to this activity, including processing, communication capacity and of course communications protocol capability. It is this latter resource, namely communications protocol addresses, that are now coming under acute pressure, as we have now managed to consume the available pool of IPv4 protocol addresses in a large part of the world.

There is a technical answer to this shortage of IPv4 addresses in the form of the IPv6 protocol that increases the supply of addresses by a massive amount. However, the cost of utilising this expanded address space is that of equipping the entire network and the entire collection of attached devices with a second protocol implementation that supports this extended address set. This new protocol offers little in the way of any other benefits: there is no improvement in the efficiency of communication, no reduced cost in the production of the devices, and no tangible early adopter advantage in deploying it. The extended inactivity on the part of the industry at large to undertake this transition to the new protocol has lead to the supply of IPv4 addresses to run out in certain parts of the Internet. As there is still an imperative for all parts of the Internet to continue to use IPv4, despite these address shortages, the industry has been forced to deploy various forms of IPv4 network middleware that allows individual addresses to be shared among multiple customers simultaneously.

The risk from continued deployment of such middleware is that the network loses flexibility and robustness, and gains fragility and imposed limitations on what applications and services can do in the Internet. The longer term outcomes of this course of action are outcomes that ultimately impair the ability of the network to support continued innovation and evolution, and this becomes a situation that favours the incumbents in both the areas of provision of carriage and the provision of content and services and implicitly creates increasing barriers to entry for new actors and innovative services, as noted in a recent study commissioned by the United Kingdom's communication regulator (OFCOM) on the implications of CGNs.³⁰

There is the prospect that the effective failure of this transition to IPv6 is now entering into a situation that increases the prospects of a market failure of the Internet economy itself. This brings renewed importance to the commitment made by Ministers in the Seoul Declaration on the Future of the Internet Economy, 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."³¹

GLOSSARY OF TERMS

CGN - Carrier Grade NAT. A NAT unit used within an ISP network infrastructure.

CIDR - Classless Inter Domain Routing. An Internet routing paradigm that passes both the network prefix and a mask of significant bits in the prefix within the routing exchange. This supercedes the earlier paradigm of classful routing, where the mask of significant bits is inferred by the value of the prefix (where Class A network prefixes infer a mask of 8 bits, Class B network prefixes infer a mask of 16 bits, and Class C network prefixes infer a mask of 24 bits). CIDR commonly is used to denote an Internet environment in which no implicit assumption exists of the Class A, B, and C network addresses. BGP (Border Gateway Protocol) version 4 is used as the de facto method of providing CIDR support in the Internet today.

DHCP - Dynamic Host Configuration Protocol. A protocol that is beginning to be used quite pervasively on end-system computers to automatically obtain an IP (Internet Protocol) host address, subnet mask, and local gateway information. A DHCP server dynamically supplies this information in response to end-system broadcast requests.

DNS - Domain Name System. The DNS is a widely distributed mapping system for the Internet that is commonly used to map names to IP addresses. This allows users to nominate network service points by a symbolic name, the DNS is capable of mapping this name into the IP address of the associated service point. The term "DNS" is often used both to describe the structure of the names used in this system, the protocol used to resolve names into mapped results, and the widely distributed network of resolvers and servers that collectively support the DNS for the Internet.

Dual-stack - A technique of equipping devices with two IP protocol stacks operating concurrently. This technique allows the device to communicate using either, or both, protocols simultaneously.

IANA - The Internet Assigned Numbers Authority. This is the entity that maintains the protocol parametre registry for the Internet's protocol suite. This functions includes the DNS name and IP address registry, as well as registries for many other protocols and functions. IANA functions operator is an activity performed by the Internet Corporation for Assigned Names and Numbers (ICANN) under terms of a contract with the United States Department of Commerce

ICANN - The Internet Corporation for Assigned Names and Numbers. This body co-ordinates the Internet Assigned Numbers Authority (IANA) functions, which are key technical services critical to the continued operations of the Internet's underlying address book, the Domain Name System (DNS). The IANA functions include: the co-ordination of the assignment of technical protocol parameters including the management of the address and routing parameter area top-level DNS domain; the administration of certain responsibilities associated with Internet DNS root zone management such as generic and country code Top-Level Domains; the allocation of Internet numbering resources; and other services. ICANN performs the IANA functions under a U.S. government contract.

IETF - Internet Engineering Task Force. An engineering and protocol standards body that develops and specifies protocols and Internet standards, generally in the network layer and above. These include routing, transport, application, and occasionally, session-layer protocols. The IETF works under the auspices of the Internet Society (ISOC).

IPv4 - Internet Protocol version 4. The version of the Internet protocol that is widely used today. This version number is encoded in the first 4 bits of the IP packet header and is used to verify that the sender, receiver, and routers all agree on the precise format of the packet and the semantics of the formatted fields.

IPv6 - Internet Protocol version 6. The version number of the IETF standardised next-generation Internet protocol (IPng) proposed as a successor to IPv4.

ISP - Internet Service Provider. A service provider that provides external transit for a client network or individual user, providing connectivity and associated services to access the Internet.

NAT - Network Address Translation. A network unit that translates the address and port fields in the IP, TCP and UDP header of packets that are passed through the unit. This unit is asymmetric with an "inside" and an "outside," and typically allows a single "outside" IP address to be shared across a number of end devices on the "inside."

RIR - Regional Internet Registry. There are at present five Regional Internet Registries. These bodies are collectively responsible for the distribution and registration of IPV4 and IPv6 protocol addresses and Autonomous System Numbers to Internet Service Providers, National Internet Registries, and various other Local Internet Registries. There are five Regional Internet Registries at present. The Regional Internet Registries also host a regional address policy community that develops address distribution policies that govern the operation of the RIR's address management functions in an open, transparent and bottom up manner. AFRINIC is the RIR that is serving the African region. APNIC is the RIR serving the Asia Pacific region. ARIN serves North America and parts of the Caribbean. LACNIC serves Latin America and the Caribbean. The RIPE Network Coordination Centre (RIPE NCC) is the RIR that serves Europe and the Middle East.

SYN scanning - A form of malicious scanning of networks, where the scanner sends packets that are identical to the opening packet of a conventional communication. For the TCP protocol this opening packet uses the "SYN" flag in the TCP options. If the packet reaches a destination device, the device may respond in some fashion, indicating that the address is being used.

TCP - Transmission Control Protocol. TCP is a reliable, connection- and byte-oriented transport layer protocol within the TCP/IP protocol suite. TCP packetizes data into segments, provides for packet sequencing, and provides end-to-end flow control. TCP is used by many of the popular application-layer protocols, such as HTTP, Telnet, and FTP.

Tunnelling - A process of encapsulation of placing an IP packet inside a IP packet header. In the context of the IPv6 transition this has been used to carry IPv6 packets across IPv4-only networks, where the IPv6 packet is "tunnelled" across the IPv4 network through the use of these added Ipv4 packet headers.

NOTES

¹ A distinction is drawn here between “end user” equipment and the entire set of equipment used in the wired Internet. The first set encompasses personal computers, laptops and other edge devices that are used as one end of the end-to-end connections made across the Internet. The larger set of equipment used by consumers includes various edge modems, routers and related equipment that is used within the network itself. The number and capability of these devices cannot be measured as readily, as they are effectively invisible to end-to-end connections. It is commonly thought that the IPv6 capability of these devices is, on the whole, far lower than that of the end user equipment. Such devices are generally not upgraded in the field, so their capabilities remain at the state when sold, and because of the price sensitive nature of this segment of the consumer market most vendors of this equipment tend to avoid adding functionality that is not an essential part of the user’s requirements www.google.com/ipv6/statistics.html.

² <http://www.oecd.org/sti/40839436.pdf>.

³ Solenksy, F. (1990), "Internet Growth," Steering Group Report, p. 61, Proceedings of the 18th IETF Meeting, August 1990. www.ietf.org/proceedings/prior29/IETF18.pdf.

⁴ Fuller, V. and T. Li (2006), "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", RFC 4632, August 2006. <http://tools.ietf.org/html/rfc4632>.

⁵ Deering, S. and R. Hinden (1995), "Internet Protocol, Version 6 (IPv6) Specification," RFC 1883, December 1995. <http://tools.ietf.org/html/rfc1883>.

⁶ Egevang, K. P. Francis (1994), "The IP Network Address Translator (NAT)", RFC 1631, May 1994. <http://tools.ietf.org/html/rfc1631>.

⁷ Metcalfe's law states that the value of a telecommunications network is proportional to the square of the number of connected users of the system. http://en.wikipedia.org/wiki/Metcalfe's_law.

⁸ This figure is based on the data published by the International Telecommunications Union, available at www.itu.int/ITU-D/ICTEYE/Indicators/Indicators.aspx.

⁹ This figure is an estimate which is based on an IPv6 end user census, undertaken by APNIC. The data is available at <http://labs.apnic.net/dists/v6dcc.html>.

¹⁰ www.verizonbusiness.com/Products/networking/internet/ipv6/.

¹¹ Hinden, R. and S. Deering (2006), “IP Version 6 Addressing Architecture”, RFC4291, February 2006. <http://tools.ietf.org/html/rfc4291>.

¹² The contract between the United States Department of Commerce and the Internet Corporation for Assigned Names and Numbers can be found at www.icann.org/en/about/agreements/iana/contract-01oct12-en.pdf.

¹³ www.nro.net/about-the-nro/regional-internet-registries.

¹⁴ www.icann.org/en/resources/policy/global-addressing/allocation-ipv4-post-exhaustion.

¹⁵ APNIC 26 Policy SIG meeting transcript, APNIC 26, Christchurch, New Zealand, August 2008. <http://archive.apnic.net/meetings/26/program/policy/transcript.html#smith-prop62>.

- 16 Final/8 APNIC announcement: www.apnic.net/publications/news/2011/final-8, RIPE NCC announcement:
 17 [www.ripe.net/internet-coordination/news/announcements/ripe-ncc-begins-to-allocate-ipv4-address-space-](http://www.ripe.net/internet-coordination/news/announcements/ripe-ncc-begins-to-allocate-ipv4-address-space-from-the-last-8)
 18 [from-the-last-8](http://www.ripe.net/internet-coordination/news/announcements/ripe-ncc-begins-to-allocate-ipv4-address-space-from-the-last-8).
 19 www.potaroo.net/ispcol/2013-08/when.html.
 20 The Internet Gateway Device Protocol (IGD),
 21 http://en.wikipedia.org/wiki/Internet_Gateway_Device_Protocol.
 22 Huston, G. (2012), Measuring Dual Stack Quality, September 2012, [www.potaroo.net/presentations/2012-](http://www.potaroo.net/presentations/2012-09-19-dual-stack-quality.pdf)
 23 [09-19-dual-stack-quality.pdf](http://www.potaroo.net/presentations/2012-09-19-dual-stack-quality.pdf).
 24 Grundemann, C. (2012) CableLabs, CGN Logging, NANOG 54, San Diego, February 2012,
 25 www.nanog.org/meetings/nanog54/presentations/Tuesday/GrundemannLT.pdf.
 26 Howard, L. (2012), Total Cost of Ownership of Carrier Grade NAT, NANOG 56, Dallas, November 2012,
 27 www.nanog.org/meetings/nanog56/presentations/Wednesday/wed.general.howard.24.pdf.
 28 www.networkworld.com/community/blog/us-government-progress-ipv6-deployment.
 29 [https://agimo2.govspace.gov.au/files/2012/04/Endorsed_Strategy_for_the_Transition_to_IPv6_for_Austral](https://agimo2.govspace.gov.au/files/2012/04/Endorsed_Strategy_for_the_Transition_to_IPv6_for_Australian_Government_agencies.pdf)
 30 [ian_Government_agencies.pdf](https://agimo2.govspace.gov.au/files/2012/04/Endorsed_Strategy_for_the_Transition_to_IPv6_for_Australian_Government_agencies.pdf).
 31 www.telecompaper.com/news/pts-says-govt-agencies-must-step-up-ipv6-deployment—941404.
www.forumstandaardisatie.nl/fileadmin/os/presentaties/10mei12_constanze-buerger.pdf.
www.v6pc.jp/en/whats/link.phtml.
www.vyncke.org/ipv6status/compare.php?metric=p&countries=us.de.au.se.jp.
 Huston, G. (2012b), "Measuring Dual-stack Quality", September 2012,
www.potaroo.net/presentations/2012-09-19-dual-stack-quality.pdf.
 Durand, A., R. Droms, J. Woodyatt and Y. Lee, (2011), "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011. <http://tools.ietf.org/html/rfc6333>.
 OFCOM (2013), "Report on the Implications of Carrier Grade Network Address Translators", OFCOM MC/159, October 2013.
http://en.wikipedia.org/wiki/Moore's_law.