Economics of IPv4 Transfer Market on IPv6 Deployment

Andrew Dul

September 2011

andrew.dul@quark.net

Table of Contents

Introduction

Internet Protocol numbers are used every day by billions of people who communicate over the Internet. These unique identifier numbers allow the computers, mobile devices, and servers on the Internet to communicate with each other. The Internet developed under a numbering system known as IPv4. The IPv4 available number pool is largely expected to be depleted in some regions starting in 2011. A new numbering scheme, known as IPv6, has been developed but has not been largely deployed. The lack of easily available IPv4 numbering resources and the lack of IPv6 compatible networks could cause a number of changes to the Internet including limiting growth, changing overall architecture, and restricting free information access. Here we examine the background of the IP addressing schemes, the economics behind the management of these scarce resources, and how these may affect the implementation of IPv6 into the Internet.

Background

Internet Protocol (IP) addresses are used as unique identifiers to connect computers on the Internet. IP addresses are often compared to telephone numbers used within the PSTN[1]. This analogy is imperfect since the telephony system relies on an underlying circuit-switched system whereas IP networks are packet switched. However, by looking at how phone number identifiers have been used over time we can attempt to draw similarities in cases where the systems have characteristics which can be compared. Similar to a phone number, IP addresses must be unique within the system that they are used.

Data is passed through the Internet in the form of packets[2]. Each of these packets contains a header which denotes the source and destination IP address for the data while in transit. The commercial Internet developed under the IPv4 (version 4) address scheme. This scheme uses a 32-bit number as the identifier and is often written in the dotted quad format (e.g. 192.168.0.1). The version 4 address scheme has just over 4 billion numbers available for allocation to endpoints and network infrastructure.

In the early 1990's the Internet Engineering Task Force (IETF) realized that eventually the scare nature of the IPv4 pool of addresses and that the available pool would be exhausted. A new protocol was developed, IPv6 (version 6). The IPv6 numbering scheme uses an 128-bit number as the identifier and has 3.4×10^{38} possible addresses. This new numbering scheme is so large exhaustion under a rational allocation scheme seems unlikely, thus IPv6 is finite but not

[1] Public Switched Telephone Network

[2] These packets are segments of data; each packet contains IP addresses which note the source and destination of the packet. These IP address headers are used by the intermediary devices, commonly known as routers, to move the packet from its source to destination.

necessarily scarce. The IETF originally planned for the IPv6 numbering scheme to be deployed alongside the currently in use IPv4 scheme in a method known as dual-stack[3]. This transition methodology was intended to allow the new IPv6 scheme to come in to popular use before the IPv4 address space was depleted. This transitional methodology did not occur for various reasons including the lack of economic or technical incentives.

History of IPv4 Allocations

Early IPv4 allocations were issued to organizations in three different classes. The class sizes of 256, 65,536, and 16,777,216 were known as class "C", class "B", and class "A" respectively. Since the differences between the class sizes are large and organizations often did not fit easily within a class, some organizations were allocated a much larger block than was required. For example: Massachusetts Institute of Technology (MIT) was allocated a class A block. This allocation methodology was inefficient in the use of numbering resources, but as an experimental network the researchers at the time did not perceive that this inefficiency would cause issues in the future.

In the early Internet data transmission architecture, these classes also had an effect on how traffic was routed through the Internet. The routing mechanisms used an address block's class as a part of the routing decision process. The routing functions were later changed with the implementation of the Classless Inter-Domain Routing[4] (CIDR) to be agnostic to class as a method of determining address block size.

Over the history of the Internet, IP addresses have been allocated to organizations under different schemes and record keeping systems. (Cannon, 2010) Today, the Internet Assigned Numbers Authority (IANA) holds the top-level records[5] in /8 blocks[6]. The Regional Internet Registries (RIRs) receive the /8 blocks from IANA and perform the functions of managing allocation the scheme, policy development, and performing the record keeping functions for their respective regions.

---

[3] In a dual-stack deployment every device has both an IPv4 and IPv6 address. A dual-stack host will attempt to use its IPv6 address first; if it is unable to make a connection using IPv6 it will use its IPv4 address to complete the communication.

[4] CIDR allowed IP addresses to be allocated and routed in blocks of varying bit-length or size which was not possible under the classfull addressing scheme. For example, under CIDR, a single block of 1,024 addresses (/20) could be allocated and routed.

[5] http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xml

[6] A /8 block is equivalent to a class A block, the "slash" notation replaced the classfull notation with the implementation of CIDR.

Today, there are five RIRs[7] each serving a specific geographic region. The RIRs are transnational non-profit member organizations with the majority of members being network operators. The RIRs operate on a cost recovery model to finance their operations. The RIR system currently uses a consensus driven stakeholder process for developing number resource policy. These policies are implemented by the RIR's professional staff. The policy development process is open to all individuals who wish to participate in the process.

RIRs issue address space to organizations on a contractual license to use basis. IP Addresses which were allocated prior to the formation of the RIRs are often called "legacy addresses." Legacy addresses were issued under similar needs based use assumptions[8], but may or may not be under formal written contracts with an RIR.

Fast forward to 2011, the IPv4 address space is now on the verge of being fully depleted. The best estimates available, predict that at least two of the five RIRs will deplete their available IPv4 address pools before the end of 2012. (Huston, 2009) (Huston, Transitional Uncertainties, 2011)

IPv6 has not been widely deployed due to a number of factors, but primarily because there was little incentive for users, network operators, and content providers to deploy IPv6. Indeed in some cases there have been disincentives for one to deploy IPv6 widely.

Allocation methods for IPv4 addresses

IPv4 addresses have been allocated on a "needs" basis. Over time the definition of "need" has changed, but fundamentally need is met when the IP numbers will be used to connect to the Internet or when interconnecting networks using the Internet Protocol system. The policy and procedures used to allocate address space has developed over time through the RIRs policy development process[9]. Today, the IP number resource policies governing the ARIN region are found in their policy manual known as the Number Resource Policy Manual[10] (NRPM).

Organizations are obliged to return unused addresses to the RIRs, but in practice this rarely occurs due to lack of incentives unless the organization ceases to exist. Under the RIRs registration services agreements[11] the RIRs do have the ability to audit organizations for usage of

[7] The American Registry for Internet Numbers, ARIN (http://www.arin.net) allocates IP address number resources in the North American region, RIPE (http://www.ripe.net) serves Europe & the Middle-East, LacNIC (http://www.lacnic.net) serves South America, AfriNIC (http://www.afrinic.net) serves Africa, and APNIC (http://www.apnic.net) serves the Asia-Pacific region.
[8] RFC 1366 - http://tools.ietf.org/html/rfc1366, RFC 1466 - http://tools.ietf.org/html/rfc1466, RFC 2050 - http://www.ietf.org/rfc/rfc2050.txt
[9] https://www.arin.net/policy/pdp.html
[10] https://www.arin.net/policy/nrpm.html
[11] https://www.arin.net/resources/agreements/rsa.pdf

blocks which have been allocated and the ability to revoke the rights to use those blocks if address policies have not been complied with, however this practice has only rarely been used. Address blocks from fraudulently obtained, bankrupt, and otherwise defunct organizations are reclaimed by the RIR and are reallocated to other organizations. Aggressive reclamation of abandoned, underutilized, or unused address resources has been debated by the RIR community for a number of years, but such an activity would be resource intensive and the potential yield for such an activity was expected to be low so these activities were not vigorously pursued. RIRs have also not engaged in reclamation activities with legacy holders of IP resources due to the lack of formal contracts between the entities. The lack of formal contracts increases the potential cost of recovery and risk due to lawsuits and other legal action when disputes would arise during the reclamation process.

IP addresses value

The IP numbers themselves do not have an intrinsic value, but the value is derived when the numbers are uniquely allocated for use in interconnecting using the Internet Protocols. This value derivation is similar to the value derived when radio spectrum is used. Spectrum itself is not valuable, but when used with equipment designed for transmitting and receiving radio signals the spectrum's use becomes valuable. Exclusivity and uniqueness are required for an IP numbering system to provide value in the same way that exclusive right to use a portion of radio spectrum can provide an organization value.

As the IPv4 address pool continued to be depleted other methods such as auctions, rationing, renewable permits, and transfer systems were considered. (Dell, 2010) Within the various regions, stakeholders debated the positive and negative aspects of changes within the registry system. (Edelmen, 2009) (Lehr, Vest, & Lear, 2008) Between 2007 and 2010 four of the five registries[12] decided to implement a system which allowed the transfer of resources between entities. These policy changes while not specifically detailing the action create a secondary transfer market for IPv4 address blocks.

The secondary IPv4 transfer market

The implementation of a directed transfer policy within the regions created a secondary market for IPv4 address blocks. From a theoretical standpoint this market should allow number resources to be transferred from entities which have an abundance of IPv4 resources to those organizations which have a need based upon economic incentives. Since it is likely that other

[12] AfriNIC, in the African region, is still considering a similar transfer policy change to allow IPv4 transfers http://www.afrinic.net/docs/policies/AFPUB-2011-v4-001-draft-01.htm

non-economic mechanisms will not redeploy IPv4 addresses in a beneficial manner to other network operators after exhaustion, a secondary market which permits the transfer of the address blocks using an economic incentive for efficient use appears to provide value to all Internet stakeholders. (Mueller, 2008)

With IPv4 exhaustion the lack of freely available additional IPv4 addresses will likely push up the cost of obtaining additional IPv4 addresses on the secondary market compared with the cost of the initial allocations. This increasing economic cost of IPv4 can provide a catalyst and economic market incentive for network operators to transition to IPv6. With a real economic cost to obtaining IPv4 addresses, the real cost of transitioning to IPv6 can now provide future value to organizations which transition.

Some stakeholders, however, have argued that the transfer market will impede IPv6 adoption, by allowing some organizations to "buy" their way out of the problem of exhaustion for a short period of time. (Doesburg, 2011) While the transfer market is likely to prolong some organizations transition to IPv6 the demand for IPv4 addresses will likely be much greater than the available supply. Delaying the deployment of IPv6 is likely to put those organizations at a competitive disadvantage compared to their competitors who adopt IPv6. **Only the transition to IPv6 has the ability to increase the supply of IP number resources by allowing some users to operate using the new IPv6 addresses.** However, as we will discuss later, IPv4 addresses are still required by IPv6 users to access the current IPv4 Internet.

In early 2011, the first public example of a large IPv4 address transfer was recorded when Microsoft paid $7.5M USD to Nortel under a bankruptcy auction for 667k IPv4 addresses. This purchase put a value on an IPv4 address of $11.25 USD on the secondary market. (Brickley, 2011) While Microsoft could have acquired IPv4 numbers from a regional registry for considerably less in economic value, it chose to purchase the rights to use the numbers from Nortel.

The creation of an IPv4 secondary market potentially provides a huge economic windfall to some organizations which either obtained abundant IPv4 "Legacy" resources or organizations which no longer have a need for the IP address resources that they hold. Based upon the known market transaction above, this values the total IPv4 address market at $43.3B USD. A number of entities, such as Hewlett Packard, received multiple large class A allocations. The equivalent two /8 allocations[13] currently held by HP have the market value of $378M USD, based upon the Nortel-Microsoft transaction.

[13] 15.0.0.0/8 & 16.0.0.0/8 - http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xml

The market price for IPv4 addresses is likely to vary over time and also vary by block size. If a large number of IPv4 addresses is offered on the transfer market this would likely push the price down in the short term, however some entities may choose to bring IPv4 supply to the market at their discretion which could result in higher overall prices if the entities were able to control a large amount of the supply of IPv4 addresses available for transfer. It is also logical to expect that the price per IPv4 address will vary by block size. Since entities with excess IPv4 resources could also control the block sizes that they offered for transfer they may be able to create additional value for their IPv4 addresses by deaggregating them into smaller blocks and offering the smaller blocks for transfer.

From a public policy perspective the trading and transfer of IPv4 resources can have a negative effect on the ability of some developing organizations and countries. The increasing cost of IPv4 resources will not assist in the promotion of the continued development of an Internet based communications infrastructure. This increased cost could further increase the economic and communication divide between the developed and under-developed countries. In response LacNIC[14] and AfriNIC[15], the registries which represent the largest number of developing countries in Latin America and Africa respectively, are developing address allocation policies to ensure that the remaining IPv4 address blocks are used for the benefit of organizations within their regions. These policies are especially relevant because it is predicted that these two RIRs will have available address blocks for some time after the other three RIRs.

The increased cost of IPv4 after all the registries have exhausted their available supply, however, may encourage IPv6 adoption in developing countries potentially putting them at a future competitive advantage. Since new IPv6 networks do not have sunk capital costs in IPv4 network equipment & resources, these organizations are not impeded by their IPv4 only infrastructure. New IPv6 networks also do not have an internal incentive to “prolong” the life of their existing infrastructure. In other industries, such as the steel industry, the adoption of newer technology has led to a long-term competitive advantage. (Crandall, 1981) The long term competitive advantage for network operators is likely to be seen in the development of human capital with the increase in technical knowledge required for IPv6 deployment and operation. The Internet itself has enabled this human capital to be widely distributed and used efficiently throughout the globe.

Since the US market is the origin of the Internet and early IP address allocations were very generous due to the classfull nature of allocations, organizations which received these initial allocations have the highest potential as a source of additional IPv4 address space after

[14] http://lacnic.net/documentos/politicas/LAC-2009-04v3-propuesta-en.pdf
[15] http://www.afrinic.net/docs/policies/AFPUB-2011-v4-003-draft-02.htm

exhaustion. This currently underutilized address space has the potential to be reallocated through the transfer market. This action would extend the life of the IPv4 or potentially provide the needed IPv4 resources necessary for the transition to IPv6. Further delays will only complicate the transition plans for organizations as they consider transition technologies and the methods for deploying IPv6 to their customers.

The transfer policies of the RIRs currently do not permit inter-regional transfers. Inter-regional transfer policies are currently being discussed within a number of regional registries, but no consensus has developed yet to create a global transfer policy for IPv4 address resources.[16] Inter-regional transfers have the possibility of both helping and hindering network expansion in developing countries. An inter-regional transfer market has the potential to bring value to the rapidly developing countries in Asia, such as China & India, by allowing the underutilized legacy resources to be transferred from North American organizations. Inter-regional transfers could also hinder development by diverting or increasing the economic costs of network deployment in countries, such as in sub-Saharan Africa, which currently do not have robust Internet infrastructure.

Demand for IPv4 addresses

The demand for IPv4 numbers under the needs based allocation scheme has been directly related to the growth in interconnected data networks and Internet users. Since IP numbers have no value outside of an interconnected network system, demand has generally grown in a manner similar to the growth in worldwide connectivity. The largest growth in IP numbers usage in the past five years has come from the Asia-Pacific region where the rapidly growing networks and economies of China and India have caused the greatest increase in demand of IP number resources. (Huston, Addressing 2010, 2011) Demand is also increasing, in all regions, from the number of Internet enabled devices such as smart mobile phones. Current consumption of IPv4 addresses exceeds 250 million addresses per year. (Huston, A Rough Guide to Address Exhaustion, 2011)

[16] http://aso.icann.org/global-policy-proposals/

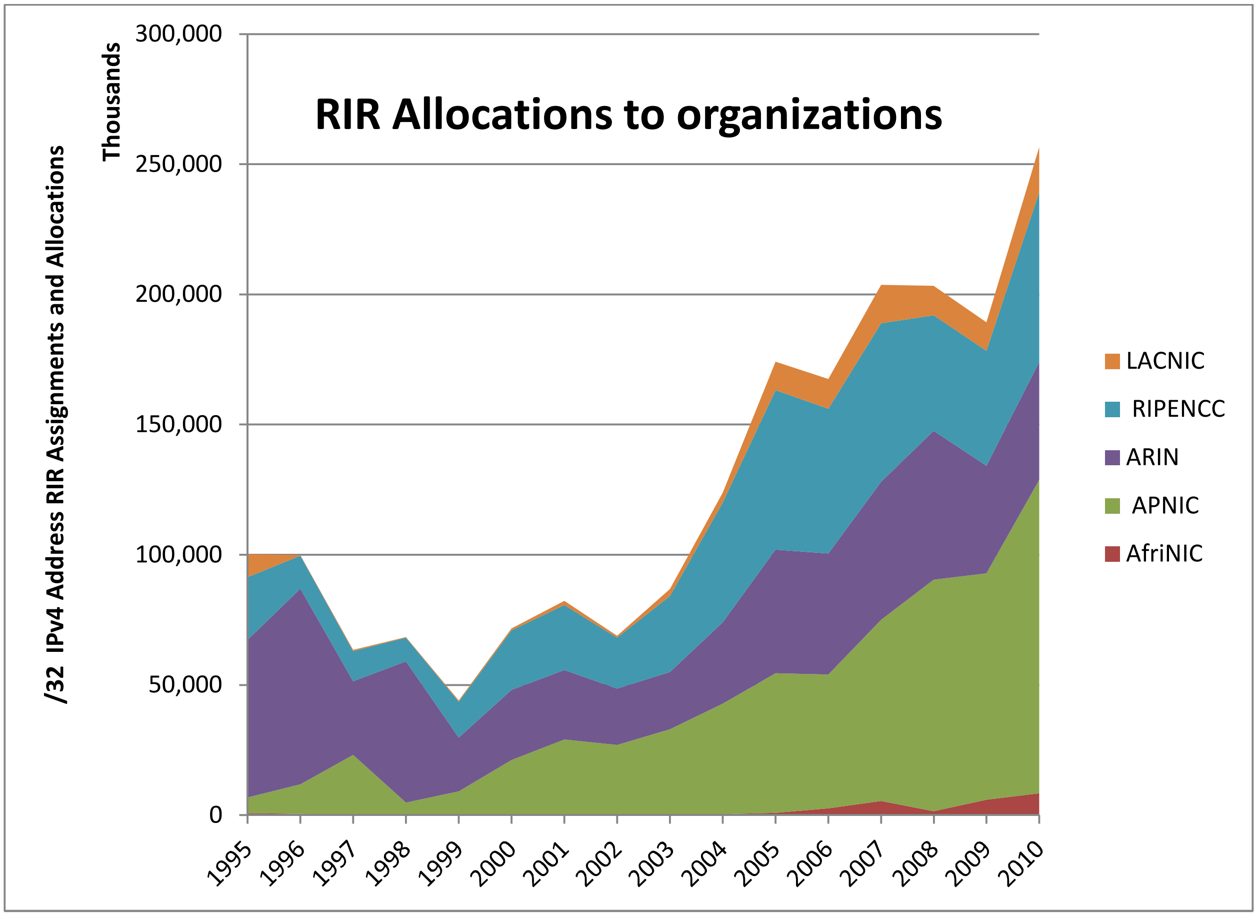

Figure 1

Figure 1 shows historical allocations by the RIRs to organizations. (Number Resource Organization, 2010) The allocations made prior to RIR formation around 1998 are considered legacy address blocks. Here you can clearly see the increase in demand from the Asia Pacific region which is managed by APNIC and the growth in the European and Middle East region managed by RIPENCC.

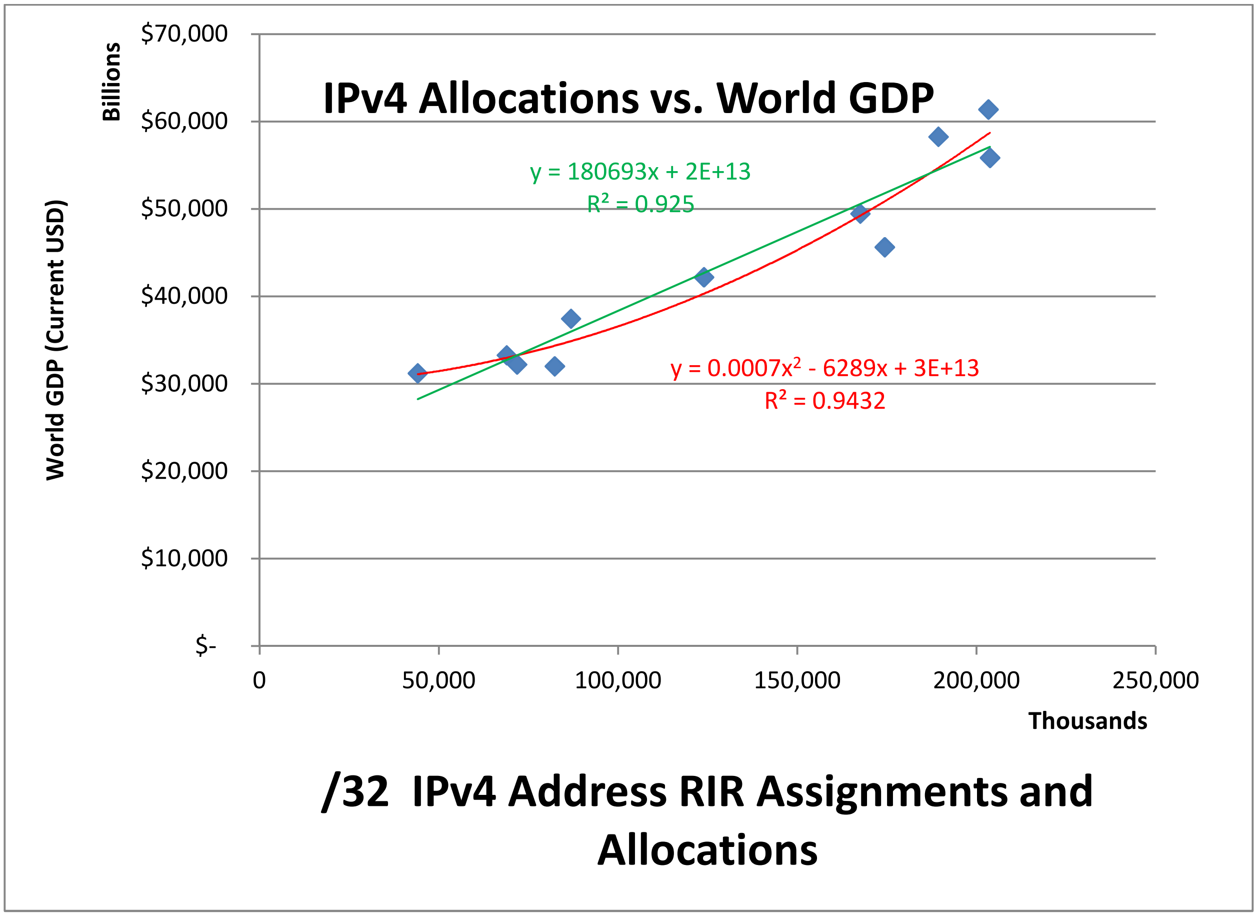

Figure 2

Figure 2 shows the RIRs yearly assignments plotted against world GDP. These data sets appear to show a strong correlation between the growth of GDP and the continued growth of the Internet. This correlation would largely be expected logically as increased world wealth has led to additional demand for technology especially Internet connected communication technologies.

IPv4 demand is not expected to diminish as IPv4 exhaustion occurs; in fact demand has increased as organization obtained additional address space in the APNIC region prior to their exhaustion of IPv4 resources in April 2011. (APNIC, 2011) It is expected that organizations in other regions may participate in a "run-on-the-bank" prior to exhaustion in each region. While IPv6 is the suggested alternative to meet the future IPv4 demands, it is not a perfect substitute. Furthermore because IPv6 endpoints cannot reach the current IPv4 Internet without transition technologies, future IPv4 demand will include the new demand for IPv4 transition technology implementations. RIRs have attempted to meet some of this new demand by reserving a block of addresses within their free pool and creating new allocation policies[17] which allow a limited number of IPv4 addresses to be allocated to an organization for transition technologies.

Debate between free transfer and needs based secondary market transfers

Three of the four[18] RIRs maintain a needs based requirement that is imposed on the secondary transfer market. Under their current policies the RIR requires the demonstration of need in order to process a transfer (registry database update) in a secondary market transaction. This needs based requirement is designed to discourage speculation in the IPv4 address market. One RIR, APNIC which represents the Asia Pacific region, has implemented a policy which does not have a needs based requirement[19] on secondary market transaction. APNIC, in its August 2011 meeting, reached consensus on a policy[20] change requiring a needs based requirement for its region, but this policy has not yet been implemented.

Some have argued that requiring a needs based review on a secondary market transfer will undermine the stability of the RIR system and specifically the registry function. (Huston, 2008) When a needs based requirement is imposed this requirement may cause some organizations to purchase address blocks on the black/grey market and not have the transfers recorded or the transfers would then be recorded in third-party non-affiliated registries. Additionally, new organizations[21] have been formed to provide 3rd party alternative registry services. The debate about the value of alternative registries has just begun. (Vixie, 2011) (Mueller, 2011) While one in general economic model terms would expect increased competition to lead to a better output, the registry function requires uniqueness and thus at its heart has a monopolistic element. Without uniqueness the value of an IPv4 address declines to the user. Non-recorded and non-coordinated 3rd party registry functions have the potential to reduce the accuracy of the RIR databases and thereby reducing the reliability and integrity of the RIR registry databases.

The RIR model is not a regulatory based system, but a unified cooperation model. The value and integrity of the RIR databases derives from the trust network operators have placed in these the RIRs to maintain an accurate accounting of how and where address blocks are allocated. If the accuracy of the primary records is undermined, the value of the database as a whole is diminished. Today, network operators use the RIR databases to ensure that organizations have been allocated specific address space before an operator will permit an organization to route specific addresses. Network operators are not under any obligation to use the RIR databases, but

---

[17] https://www.arin.net/policy/proposals/2008_5.html
[18] AfriNIC does not currently have a directed transfer policy.
[19] http://www.apnic.net/policy/proposals/prop-050
[20] http://www.apnic.net/policy/proposals/prop-096
[21] http://www.depository.net

the RIR's stewardship of the Internet Protocol records history and its current vetting of new allocations provide value to network operators in their fight against unsolicited email and other fraud that is perpetrated on the Internet.

Since the increase in demand for IPv4 numbers in a non-speculative market are tied only to the growth in the Internet endpoints and associated infrastructure; any demand increases in a speculative market would likely to cause prices for IPv4 number resources to rise further as a rational speculator would require a return on their investment of address blocks. Based upon rough growth projections and potentially available IPv4 blocks it is estimated that IPv4 addresses which may become available as a result of the transfer process would only likely meet the non-speculative demand for an additional 1-2 years.

In order to prevent speculation, the stakeholders within the RIR community created policies which restricted the transfers between organizations to those entities which could show that they have a justified need for the number resources. Some policies[22] also have a policy time limit which only allows a certain number of transactions within a specific period to further reduce the likelihood of rampant speculation.

The routing table

Packet data traffic is moved through the Internet through the use of intermediary systems, commonly known as routers.[23] The routing table is created using a protocol known as Border Gateway Protocol (BGP). BGP allows routers to dynamically and automatically communicate destination information to each other. Each router announces to its neighbors what IP address blocks it can reach. Collectively these network announcements are correlated into a routing table. BGP allows network operators to control traffic in and through their network forming the administrative domains of the Internet. An administrative domain is defined by an autonomous system number[24] (ASN). Each routing table entry contains a list of ASNs[25] which denote the path that the route announcement has made through the network. These announcements are the data necessary for the traffic management command and control system which directs traffic throughout the Internet.

---

[22] RIPE's policy limits each member to one transfer every 2 years, LACNIC policy limit is 1 year.

[23] Every router along the path from source to destination examines the header of the IP packet which contains the source and destination IP address. To determine where a router should send the packet it consults an internal routing table which lists the destination IP address blocks and the next-hop (intermediary system). By matching the destination IP address from a packet and the information contained in the routing table the router can determine where to send the packet.

[24] http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_9-1/autonomous_system_numbers.html

[25] Autonomous system numbers are also distributed by the RIRs. These numbers are assigned to network operators and are a requirement of the BGP protocol.

Each new IP address block that is added to the Internet creates an additional routing table entry on each of the intermediary routers[26]. Over the development of the Internet, the routing table size has been issue which caused concern for network operators. Each routing table entry takes up memory on the router. If the routing table grows beyond the physical capacity of the network hardware the router could cease functioning creating a network outage. Limitations in routing table scaling extend beyond the memory required to hold the routing table entries.[27] With a larger number of entries there is a delay in the ability for the network to fully converge, that is for the network to be able to pick a stable usable traffic path through the Internet. Hardware[28] also must be scaled to allow for additional routing table entries, some routers have special fast lookup tables known as TCAMS which have significant power requirements which may not scale to being able to handle millions of routes.[29]

Under the original classfull model a class C network was often too small for most operators, a class A was too large, and a class B was often a good size. There were however only a 16,384 class B entries available in the original classfull addressing scheme. Over time, the number of class B entries began to dwindle and registry started assigning large numbers of class C block instead of a class B block. This change in the classfull addressing model did not efficiently use routing entries. The smallest class C network entries which required one routing table entry for 256 IPv4 addresses created additional routing table entries than was necessary[30]. This inefficiency was one of the driving forces behind the creation of CIDR. With the implementation of CIDR, network entries for smaller blocks could be combined and this reduced the growth of the network entries in the routing tables as smaller blocks could be combined.

While the CIDR system was more efficient in its use of routing table entries, large CIDR blocks could still be broken into multiple pieces and used separately. When an address block is broken into multiple pieces and routed separately a block is said to be deaggregated. Intentional deaggregation is done for a number of reasons including moving of networks to new locations, assigning networks to other organizations or customers, traffic-engineering, security, and the practice of using multiple network operators for redundancy (multi-homing). Deaggregation can also be unintentional and is often results from misconfiguration of network devices and misunderstanding of IP protocols including BGP and CIDR.

---

[26] These routers are also commonly called "default-free zone" routers.
[27] https://www.arin.net/participate/meetings/reports/ARIN_XX/PDF/wednesday/RoutingTable_Schiller.pdf
[28] http://www.nanog.org/meetings/nanog21/presentations/li.ppt
[29] RFC 4984 - http://tools.ietf.org/html/rfc4984
[30] For example, an organization with 1000 end-points would require 4 class C networks and 4 network entries under the classfull addressing scheme, but under CIDR addressing 1000 end-points could be accommodated with one /20 address block and one routing table entry.

While deaggregation[31] creates additional routing table entries, it provides end-users and network operators with a valuable mechanism for controlling traffic to and from IP addresses. Every deaggregation causes additional routing table entries to be created, while this action provides value to the organization which performed the deaggregation, the cost of this action is borne by all other network operators who must provision their network routers to accept and use these additional deaggregated route table entries.

Exchanges of routing information are often done through exchange points or private interconnections. The routing system, however, does not directly limit or provide an economic incentive to limit the number of routes which can be announced and thus a single actor can have a large impact on all other network operators. This action creates an externality which is not mitigated by the current routing exchange systems. (Mueller, 2010)

The physical limits of router hardware have forced network operators to work to limit routing table growth and this has been done by various non-economic practices. These practices include the IP address allocation policy, public "shaming"[32] of network operators who do not follow best practices, and "cooperative norms"[33] which have developed to limit growth.

Today, the routing table is composed of more than 350,000 network entries. Additional growth is expected by network operators, but a large increase could create instability in the Internet. Moving suddenly to a system of millions of entries would not be economically and physically possible with today's hardware and software. Physical hardware limits have also often defined the maximum number of network entries available within a router, while those physical limits have expanded with the previous growth of the routing table, large changes may not be possible with technology available today. Because of the limitations of routing table growth on router hardware, the Internet community has used IP address allocation policies as a method to control the growth of network entries.

Network operators in general, have an internal economic incentive to limit the growth of the routing table. A smaller routing table leads to lower capital and operating expenses and prolongs the life of existing hardware. Operators have sought to cooperate without collusion to limit growth of the table.

The creation of a secondary market for IPv4 network addresses is expected to cause a measureable increase on the number of routing table entries. These entries would be created as

[31] Deaggregation by end-users for redundancy and traffic control is commonly called multi-homing and deaggregation by providers for traffic control is commonly called traffic-engineering.
[32] CIDR Report - http://www.cidr-report.org/as2.0/
[33] The best example of these cooperative norms are the use of routing filters which limit the acceptance of routing entries which are smaller than a /24 or 256 IPv4 addresses.

organizations with large unused blocks deaggregate the blocks and transfer them to smaller organizations. Further deaggregation is expected as organizations use techniques such as NAT to extend their IPv4 address resources, possibly between multiple locations. While some organizations will need very large blocks of address space it is hypothesized that the smaller blocks may produce higher transfer fees per IP address to the selling entity since existing large organizations will either pursue IPv6 deployment or become more efficient with their existing IPv4 address resources. Each smaller block which is created by the transfer market through deaggregation creates an additional routing table entry.

Technical methods (Scudder, 2007), such as Locator Identifier Separation Protocol (LISP) (Meyer, 2008), are being created to separate the routing function from the current inherent link with address resources. This technology has not yet gained broad acceptance and the economic costs of switching to the new routing system are high and require multiple network operators within the industry to cooperate.

IPv6 deployment incentives (or lack thereof)

Some have argued that IPv6 has not been widely deployed and adopted because there have not been economic incentives for organizations to move to the new protocol. Indeed, implementations of IPv6 require additional capital and operating expenses. Additional training and configurations are required to operate an IPv6 network and depending on the age of an organization's Internet infrastructure moving to IPv6 may require new hardware or software. While large networking vendors such as Cisco, Juniper, Microsoft, and Apple have implemented IPv6 in most of their products, production support and reliability are less than their IPv4 counterparts. Additionally many small software and hardware developers do not support IPv6 because they did not have any customer requirements or near-term economic incentive to support the new protocol.

When looking at exhaustible resources, a transition to another technology, in this case IPv6, will not occur until the price of the current resource (IPv4) exceeds the cost of the new resource (IPv6). (Elmore, Camp, & Stephens, 2008) Since IPv6 technology has been more costly to deploy compared with IPv4 the "Hotelling Rule" applies and IPv6 transition will not occur until the IPv4 resources are exhausted. (Hotelling, 1931)

Furthermore, some smaller and midsize organizations may currently have adequate IPv4 resources to serve them in the near future. Until IPv6 has a significant deployment penetration, provides additional benefits, has a regulatory mandate, or has the stability and reliability of IPv4 these organizations will have little or no economic incentive to adopt IPv6. Additionally since initial development, deployment, and operations costs for IPv6 will be high during the initial phases, these well provisioned IPv4 organizations have an economic incentive to wait for costs of deploying IPv6 to drop as expertise and experience with the protocol becomes more prevalent in the marketplace.

New and high growth organizations are those organizations which will have economic incentives to deploy IPv6 after IPv4 exhaustion. The inability to obtain IPv4 IP addresses or the cost of obtaining IPv4 addresses on the secondary market will create an economic incentive for these organizations to consider IPv6 network deployment.

IPv4 address substitutes

With the depletion of IPv4 address resources, it is anticipated that IPv6 networks will start to be deployed. However, IPv4 addresses and IPv6 addresses are not perfect substitutes. That is the value of an IPv4 address is greater than the value of an IPv6 address within the current Internet architecture. The difference in value between the number resources is derived from the fact that the newer IPv6 addresses are not backward compatible with IPv4 numbers.

An endpoint with an IPv6 address cannot directly reach resources on the current IPv4 Internet. For an IPv6 only endpoint to reach the current IPv4 network a transition technology must be used. A number of transition technologies have been created that allow an IPv6 endpoint to reach an IPv4 resource, but these transition technologies require the use of IPv4 addresses. (Huston, Transitioning Protocols, 2011) These transition technologies are also technically complex and have been shown to have a negative impact on data transaction performance compared with native transactions. (Huston, Stacking it up, 2011)

These transition technologies create an additional demand for IPv4 resources. Large organizations which transition their endpoints to IPv6 may be able to reclaim IPv4 resources from these endpoints to use with transition technologies to offset this demand.

Network Address Translation[34] (NAT) is a technical method which allows multiple devices to use a single IP address. The most common use of this technology is in the home gateway routers which are found in many homes which are connected with broadband Internet service. This technology is also widely used in corporations as a security mechanism. While this technology allows multiple devices to use a single IP address, the technology prohibits certain types of data communications. The limitation of this technology makes NAT a poor substitute for a native IP address for many applications.

[34] http://en.wikipedia.org/wiki/Network_address_translation

Network address translation can be considered a close substitute for additional IPv4 addresses when the number of IPv4 addresses is limited. Indeed, NAT's popularity within the home residential market grew largely because broadband network operators did not assign additional IPv4 addresses to home customers as customers increased the number of IP enabled devices even though the technology was available to allow such assignments and such assignments were permitted under IPv4 address allocation policies. Instead network operators discouraged customers from using additional IPv4 addresses by creating pricing policies which charged customers for additional IPv4 addresses. Customers opted for a single capital purchase of a "home gateway" router running NAT instead of a recurring charge for IPv4 addresses from their network operator. Furthermore, these home gateways also provided valuable features to home users such as 802.11 wireless access points which were not offered on network operator provided customer premise equipment.

The increase in the use of network address translation within the core of the Internet has the ability to extend the life of the existing IPv4 addressing architecture. Such an extension, known as carrier grade NAT or NAT444, would likely create additional technical complexity, increase operator's capital and operational expenditures significantly for NAT hardware, and reduce customer's functionality by technically limiting available data transaction functions.

IPv6 transition, a coordination game?

Using the US market as an example, the vast majority of large IP addresses blocks are used by a few (<10) organizations. These organizations which are composed of national telephone companies (AT&T & Verizon), the nationwide cable companies (Comcast & TimeWarner), and other mobile phone and broadband network providers. Since IPv6 dual-stack adoption has not occurred before IPv4 exhaustion, these companies are now facing large scale choices regarding transition mechanisms to support their businesses with increased growth. (Gallaher & Rowe, 2006)

Early in the IPv6 development cycle, it was assumed that these large network operators would act rationally and adopt IPv6 to avoid the current exhaustion issue. The adoption of IPv6 by a single large organization would drive others to compete on technology rather than economic aspects alone. Unfortunately IPv6 technology has not become a product differentiator and IPv4 technology has continued to dominate the Internet. Some of these effects are also likely due to the sunk costs that the network operators had within their current networks. Their operational practices, network infrastructure, and customer premise equipment all easily supported IPv4 and the supplies of IPv4 were plentiful. The technology incentives were not being driven by customers and even technology vendors who had an economic incentive when network hardware was being replaced did not seize the opportunity to strongly promote IPv6. This may be a case where maximum short-term present value trumped potential long-term future value.

Since there could be competitive advantages or disadvantages for choosing one technology over another network operators may now be facing a coordination game dilemma. The large companies have largely not announced their transition and IPv6 deployment plans and thus the first mover may have an advantage in setting of the transition direction. Other markets, such as Australia, where market power is concentrated in two large entities (Telstra & Optus) appear to face similar issues. Whether these large providers are able to reach an equilibrium on transition technologies or transition plans is yet to be determined.

The lack of a clear leader in technology or by a large network operator appears to have led to a type of stalemate that has prevented the industry from easily adopting a transition methodology. The high levels of uncertainty and lack of clear path forward have caused organizations to delay future technology adoption.

IPv4 Address Market & IPv6 Adoption Economic Hypotheses

The IP address markets are complex and have been very lightly studied. Since the initial allocations were done through a needs based mechanism without specific economic market incentives, the new market incentives available in the secondary IPv4 transfer market could have a significant impact on the IPv4 address market and the Internet as a whole.

Here we look at a few hypotheses which could be tested, these hypotheses were formed based upon economic models, known factors about the industry, and estimated potential market.

Fundamentally a hypothesis is an unproven statement, the hypotheses here are postulated based upon the current state of the industry, changes in the underlying infrastructure, regulation, or other factors could cause changes in the market. Also, these hypotheses may not accurately describe the market and could be false. The hypothesis here and its description is followed by a proposed method to validate or invalidate the hypothesis.

Hypotheses #1

Large scale adoption of IPv6 will not occur until IPv4 exhaustion is complete.

As of September 2011, IPv4 exhaustion has occurred within the Asia Pacific Region when APNIC announced[35] that it had exhausted its free pool of available IPv4 addresses on April 15, 2011. The other four regional registries still maintain available IPv4 address space. The North

[35] http://www.apnic.net/publications/news/2011/final-8

America registry (ARIN) and the European-Middle East registry (RIPE) are expected to exhaust their free pools of IPv4 addresses in 2012-2013.

Research Strategy

In some ways, this hypothesis might seem inevitable since IPv6 adoption has not yet occurred and exhaustion will occur shortly. However, it is possible that IPv6 is not adopted and other technologies emerge as a replacement or IPv4 extensions such as NAT are used instead of IPv6.

Traffic measurements taken by a number of research organizations[36] show the amount of IPv6 traffic[37] on the Internet. In 2008 these measurements show that only 0.002% of all Internet traffic was IPv6. (Ringberg, 2008) If IPv6 is adopted these traffic measurements[38] will show a sharp increase in IPv6 traffic these changes can be correlated with the exhaustion dates.

Hypotheses #2

Transfer prices in various regions will be different due to the different rules, different supply, and different demand. The transfer price for IPv4 addresses via the APNIC registry will be higher on average per IPv4 address compared with transfers which occur in other regions.

The APNIC region does not currently have a needs based requirement for transfers and thus any entity could use the transfer policy to obtain IPv4 address space. This allows for speculators to obtain address space and then resell the space for a profit. The Asia-Pacific region is also the fastest growing region and shows the highest demand for IPv4 addresses. The needs based requirements in other regions will prevent some speculative buyers from obtaining IPv4 addresses in those regions as long as the needs based policies remain in effect. The North American region also has the largest amount of legacy address space. This address space seems most likely to be offered in the transfer market. The increased supply from the legacy address space is likely to push down the price in the ARIN region. The lack of inter-region transfer policy also largely prevents white market transfers between regions from harmonizing the transfer price between regions.

Research Strategy

While it is likely that a large majority of transfer transactions will occur between parties without public disclosure, regulatory and other requirements will also likely produce evidence of IPv4 transfer market transactions.

---

[36] http://blog.caida.org/best_available_data/2011/04/28/data-on-current-status-of-ipv6-deployment/
[37] http://blog.caida.org/best_available_data/2011/04/29/caidas-ipv6-measurement-and-analysis-activities/
[38] http://www.arbornetworks.com/ipv6-report-network-security-research.html

The Nortel & Microsoft transaction provided evidence of the first large transfer transaction. Other transactions are likely to be reported in the media and bid/offer prices are likely to become available as the number of transactions increases. Sales of IP address blocks have been listed on ebay[39] for a number of years and are usually removed prior to completion of the auction. It is possible that transactions will be allowed to be completed on ebay or other auction sites in the future as the transfer market becomes better understood. Availability, demand, and offering prices may also be disclosed through listing services. ARIN has setup a listing service[40] for its region which is intended to facilitate transfers between entities.

Hypotheses #3

The rational price for IPv4 addresses will be driven by the underlying revenue which can be extracted from the use of the resource. IPv4 resources will be used in locations where an organization can obtain additional revenue and normal lower margin consumers will be migrated to IPv6 as it becomes available.

Large network operators will use the exhaustion of IPv4 as a method to extract additional revenue from high-end customers. These customers (especially business customers) will be willing to pay for the known reliability of IPv4 native addresses. Low margin customers will likely be the first to be migrated to IPv6 with a transition technology. This is similar to the now prevalent use of NAT gateways in the residential broadband market. NAT gateways were adopted because NAT was "good enough" and network operators were not willing to assign multiple IPv4 addresses to residential customers without additional revenue.

Research Strategy

Pricing data for IPv4 static addresses, such as the data in Table 3, is widely available from a number of network operators. This data could be used to calculate the present value of the revenue streams associated with offering these additional services. Using this and other data one can create pricing models for IPv4 addresses under certain conditions. There are a number of ways that an IP address can be used and how they are used also has an effect on their value. There are a number of cases where IPv4 addresses are required and these functions would place the highest value on having an IPv4 address. Two common locations where IPv4 addresses are required is in the use of IPv4 static addresses on broadband Internet connections and static IP assignments with web hosting services. IPv4 static addresses are also required for a number of technology applications such as secured encrypted HTTPS[41] web transactions.

39 http://cgi.ebay.co.uk/IPv4-Address-Block-Class-B-65-536-unused-IP-Addresses-/200594720065
40 https://www.arin.net/resources/request/stls.html
41 http://en.wikipedia.org/wiki/HTTP_Secure

Using an average of the five reference values listed in Table 3 we find a mean value of $9.20 USD per month in recurring revenue. Assuming a required rate of return of 6%, a life-span of 8 years, and no salvage value, a basic cash flow analysis would produce a present value for an IPv4 address of the following for static IP for broadband & webhosting.

	Average	Lowest	Highest
I/Y	6%	6%	6%
N	8	8	8
FV	0	0	0
PMT	$110.40 = ($9.10 * 12)	$24.00 = ($2.00 * 12)	$179.40 = ($14.95 *12)
PV	$713.52	$149.03	$1114.03

This calculation produces a present value of an IPv4 address with a range of $149 - $1114. All of these values are much greater than the recent Nortel/Microsoft market transaction. Indeed they represent a 7x to 10x return based upon this analysis.

It is likely however that this simple analysis does not represent the true value of an IPv4 address. The products offered under these prices are services which require and IPv4 address and have no available substitute. One would expect that other uses of IPv4 addresses would have a much lower value.

Using the similar assumptions and the Microsoft/Nortel transaction price we can compute an estimated monthly cash flow if Microsoft used these IPv4 addresses in a similar manner.

I/Y	6%	6%	8%
N	8	4	8
FV	0	0	0
PV	$11.25	$11.25	$11.25
PMT	$1.81	$3.25	$1.95

Here we see that Microsoft's costs on an ongoing basis for similar use would require revenue of $1.81 to $3.25 to cover the cost of acquisition of these with some variation in required return and duration. With the estimated life of 8 years and return of 6% we see that using these resources in a broadband static IP or webhosting context produces a margin in excess of 500%.

Conclusions

The IP addressing environment is a complex mix between technical, economic, and other drivers. The number of stakeholders is quite varied and the motivational incentives between the stakeholders also can be quite different. With the exhaustion of IPv4 resources the Internet community is entering a new era of scarcity of numbering resources. IPv6 deployment using dual-stack did not proceed as was originally envisioned by the IETF and the current situation will likely require a number of transition technologies to allow the Internet to continue to grow. To meet future addressing and identifier needs IPv6 appears to be the only viable option assuming the continuation of the existing Internet architecture.

The creation of an IPv4 transfer market has the ability to introduce an economic incentive to the allocation scheme. This incentive can be used to encourage IPv6 adoption but could also delay adoption. A rational stakeholder, however, would have to assume that a delay in adoption while undesirable from a short-term perspective may indeed ease the transition from IPv4 to IPv6. Given the lack of IPv6 deployment today, any additional resources that can be used to smooth the transition, while still encouraging the transition, will provide stability to the Internet as a whole.

Demand for IP addresses is increasing and as we have seen is tied to the growth of the world's economies. Ensuring an adequate supply of number resources will allow the Internet to continue to bring information to all the peoples of the world.

Given that the IPv4 transfer market has been established and the first transactions have occurred, it makes sense to consider the mechanisms that are in place within the IPv4 transfer market. The current transfer policies have mechanisms to limit speculation and these seem rational, but if they undermine the allocation system this could have a negative effect on the Internet over the long term.

While today a native IPv4 address does not have a perfect substitute. The large scale implementation of IPv6 and IPv6 transition technologies will bring an IPv6 address more in line as a substitute for IPv4. Investment in IPv6 and deployment of IPv6 enabled networks will reduce the elasticity between IPv4 & IPv6 addresses as substitutes. NAT can provide a short-term substitute for additional IPv4 address and a bridge to IPv6; however the prevalent use of service provider or carrier-grade NAT within the network infrastructure will fundamentally change the architecture on which the Internet was founded. Large scale NAT deployments are also likely to be technically complex, operationally expensive, and capital intensive. Long-term development of networks based upon service provider NAT seems undesirable due to limited scalability.

The routing table growth has been a cloud that has hung over the RIR stakeholders for decades. While stakeholders often state that the RIRs do not make routing policy, their actions certainly form routing policy indirectly. Network operators are free to ignore the RIR actions, but the RIR's leadership and the actions of the RIR stakeholders individually provides needed coordination to network operators for setting of operational norms for routing. While continued growth of the routing tables can be tolerated, large scale changes without offsetting changes will likely reduce overall stability of the routing system and the Internet as a whole.

Finally, since the IPv6 transition does not have any clear leaders or regulatory requirements, the current state of operations could be considered a game of coordination or a game of chicken. Being the first to market would have advantages and/or disadvantages. World IPv6 day[42] is loosely coordinated attempt to break the current status quo. With multiple operators and content providers agreeing to test IPv6 on June 8, 2011, the negative downsides to a single organization can be limited.

[42] http://www.worldipv6day.org/

References

APNIC. (2011, April 11). *APNIC Update.* Retrieved May 25, 2011, from APNIC: https://www.arin.net/participate/meetings/reports/ARIN_XXVII/PDF/Monday/huston_apnic_update.pdf

Brickley, P. (2011, March 23). Nortel selling Internet addresses to Microsoft for $7.5 million. *Dow Jones Daily Bankruptcy Review* .

Cannon, R. (2010, December). *Potential Impacts on Communications: From IPv4 Exhaustion & IPv6 Transition.* Retrieved January 5, 2011, from Federal Communications Commission: http://www.fcc.gov/working-papers/potential-impacts-communications-ipv4-exhaustion-ipv6-transition

Crandall, R. W. (1981). *The US Steel Industry in Recurrent Crisis.* Washington, DC: The Brookings Institution.

Dell, P. (2010). Two economic perspectives on the IPv6 transition. *Info, Vol. 12 Iss: 4* , 3-14.

Doesburg, A. (2011, January 24). *Upgrade opens door to virtually infinite internet*. Retrieved January 24, 2011, from New Zealand Herald: http://www.nzherald.co.nz/news/print.cfm?objectid=10700892

Edelmen, B. (2009). Running Out of Numbers? The Impending Scarcity of IP Addresses and What To Do About It. *Harvard Business School NOM Unit Working Paper No. 09-091* .

Elmore, H., Camp, L. J., & Stephens, B. (2008, June). *Diffusion and Adoption of IPv6 in the ARIN Region.* Retrieved December 29, 2010, from Workshop on the Economics of Information Security: http://weis2008.econinfosec.org/papers/Elmore.pdf

Gallaher, M. P., & Rowe, B. R. (2006). The Costs and Benefits of Transferring Technology Infrastructures Underlying Complex Standards: The Case of IPv6. *Journal of Technology Transfer, 31* , 519–544.

Hotelling, H. (1931). The Economics of Exhaustible Resources. *The Journal of Political Economy* , 137-175.

Huston, G. (2011). A Rough Guide to Address Exhaustion. *Internet Protocol Journal, Vol. 14: Num. 1* , 2-11.

Huston, G. (2011, January). *Addressing 2010.* Retrieved January 7, 2011, from The ISP Column: http://www.potaroo.net/ispcol/2011-01/addresses-2010.html

Huston, G. (2009, May). *Predicting the End of the World.* Retrieved May 2011, from The ISP Column: http://www.potaroo.net/ispcol/2009-05/ipv4model.pdf

Huston, G. (2011, April). *Stacking it up.* Retrieved May 23, 2011, from American Registry of Internet Numbers: https://www.arin.net/participate/meetings/reports/ARIN_XXVII/PDF/Tuesday/huston-dualstack.pdf

Huston, G. (2008, November). *The Changing Foundation of the Internet: Address Transfers and Markets.* Retrieved May 25, 2011, from The ISP Column: http://www.potaroo.net/ispcol/2008-11/transfers.pdf

Huston, G. (2011, September). *Transitional Uncertainties.* Retrieved September 20, 2011, from The ISP Column: http://www.potaroo.net/ispcol/2011-09/exhaustion.pdf

Huston, G. (2011). Transitioning Protocols. *Internet Protocol Journal, Vol. 14: Num. 1* , 22-45.

Lehr, W., Vest, T., & Lear, E. (2008). Running on Empty: the challenge of managing Internet addresses. *36th Research Conference on Communication, Information, and Internet Policy*. George Mason University, Arlington, VA.

Meyer, D. (2008). The Locator Identifier Separation Protocol (LISP). *Internet Protocol Journal, Vol. 11: Num. 1* , 23-36.

Mueller, M. (2011, August 15). *ARIN and Vixie get nervous about competition.* Retrieved 15 2011, September, from Internet Governance Project: http://blog.internetgovernance.org/blog/_archives/2011/8/15/4877516.html

Mueller, M. (2010). Critical resource: An institutional economics of the Internet addressing-routing space. *Telecomunications Policy* , 405-416.

Mueller, M. (2008, July 20). *Scarcity in IP addresses: IPv4 Address Transfer Markets and the Regional Internet Address Registries.* Retrieved December 29, 2010, from Internet Governance Project: http://www.internetgovernance.org/pdf/IPAddress_TransferMarkets.pdf

Number Resource Organization. (2010). Internet Number Resource Status Report. *American Registry of Internet Numbers*, (p. 4). San Juan, Puerto Rico.

Ringberg, H. (2008, July). *A One Year Study of Internet IPv6 Traffic.* Retrieved June 1, 2011, from Nanog: http://www.nanog.org/meetings/nanog44/presentations/Tuesday/Ringberg_measurement_N44.pdf

Scudder, J. (2007, October 17). *Routing/Addressing Problem Solution Space.* Retrieved September 15, 2011, from ARIN XX: https://www.arin.net/participate/meetings/reports/ARIN_XX/PDF/wednesday/SolutionSpace_Scudder.pdf

Vixie, P. (2011, July 20). *Arrogance in Business Planning.* Retrieved September 15, 2011, from Association for Computing Machinery: http://queue.acm.org/detail.cfm?id=2008216

About the author:

Andrew Dul is currently a networking consultant at Cascadeo Corporation and has been working in various Internet networking positions since 1996. He moved to Seattle in 1998 to join the start-up network service provider, Internap. Following Internap, Andrew went to work on the Internet in the sky at Connexion by Boeing. During his time at The Boeing Company he helped develop the global network that supported the first commercial inflight Internet service. He also helped engineer the first IP based GSM & CDMA flying pico cell demonstration flights in 2005. After Boeing, Andrew served a Seattle based nationwide law firm as their Systems and Network Architect before joining Cascadeo in 2010. Andrew has been involved with the American Registry for Internet Numbers (ARIN) for over 10 years helping to contribute to the development of global IP number resource policy. Andrew holds a bachelors of science in Electrical Engineering from the University of California, Davis and is currently enrolled in the Masters in Business Administration program at Seattle University, Albers School of Business and Economics.

Acknowledgements:

I wish to thank all those who encouraged me to write this paper and to continue to think about how I could meld my Internet work and my graduate studies at Seattle University. Special thanks to Brian Kelly at Seattle University for being open to discussing this crazy idea for a paper, Geoff Huston for his insight, and to Cathy Aronson and others who reviewed drafts of this paper.

Appendixes & Data

RIR Allocations to Organizations

Year	AfriNIC	APNIC	ARIN	RIPENCC	LACNIC	Total
1995	866304	6043136	60459008	24058448	8902912	100329808
1996	547584	11335680	75181824	12611320	342272	100018680
1997	49152	23210240	28206080	11651544	468992	63586008
1998	114688	4773376	54244864	9063648	290816	68487392
1999	49152	9178624	20661760	13772416	442368	44104320
2000	516096	20759552	26866432	22911360	768512	71821952
2001	354304	28726272	26756408	24908800	1585920	82331960
2002	198144	26895360	21597696	19582728	643072	68917000
2003	210432	32904448	21915848	29231968	2603520	86866216
2004	482048	42474496	31152896	46059200	3798784	123967424
2005	937984	53633792	47431424	61323728	10941440	174399440
2006	2672128	51407360	46549504	55529608	11420160	167578760
2007	5530880	69608704	53030912	60844192	14730752	203745440
2008	1579776	88868096	57173760	44395504	11314176	203331312
2009	5991424	86976000	41291008	44174608	10934016	189367568
2010	8520960	120384000	45239808	65135968	17278976	256559712
2011	2807296	104631296	13628928	21564656	7957760	150589936

Table 1

Special Thanks to Geoff Huston of APNIC for providing this data.

World GDP by Year in Current USD

Year	World GDP (current US$)
1995	$ 29,692,820,183,841
1996	$ 30,313,420,349,692
1997	$ 30,214,893,326,118
1998	$ 30,076,187,744,326
1999	$ 31,204,194,358,536
2000	$ 32,209,707,979,350
2001	$ 32,008,721,297,934
2002	$ 33,273,921,991,935
2003	$ 37,447,356,905,510
2004	$ 42,196,337,997,515
2005	$ 45,630,781,401,164
2006	$ 49,459,976,902,212
2007	$ 55,853,287,909,433
2008	$ 61,379,607,590,518
2009	$ 58,259,785,029,004

Table 2

http://databank.worldbank.org/ddp/home.do

Static IPv4 Address Cost

Company	Cost per month for 1 IPv4 address (USD)
Qwest[43]	$5.95
O2 (UK)[44]	$8.15 (£ 5.00)
Bluehost[45]	$2.00
BellSouth (AT&T)[46]	$14.95
Comcast Business[47]	$14.95

Table 3

[43] http://www.qwest.net/help/static_ips.html
[44] http://www.broadbandchoices.co.uk/why-get-a-static-ip.html
[45] https://my.bluehost.com/cgi/help/541
[46] http://www.bellsouth.com/consumer/inetsrvcs/inetsrvcs_fa_static_ip.html
[47] http://business.comcast.com/internet/ipaddress.aspx

Current RIR Transfer Policies (June 2011)

APNIC

3. Transfers of IPv4 addresses

APNIC will process and record IPv4 address transfer requests between current APNIC account holders subject to the following conditions.

3.1 Conditions on the IPv4 address block

- *The minimum transfer size is a /24.*
- *The address block must be:*
- *In the range of addresses administered by APNIC*
- *Allocated or assigned to a current APNIC account holder*
- *The address block will be subject to all current APNIC policies from the time of transfer.*

3.2 Conditions on source of the transfer

The source entity:

- *Must be a current APNIC account holder*
- *Must be the currently registered holder of the IPv4 address resources, and not be involved in any dispute as to the status of those resources*
- *Will be ineligible to receive any further IPv4 address allocations or assignments from APNIC for a period of 12 months after the transfer, or until the exhaustion of APNIC's IPv4 space (that is, until the commencement of the use of the "final /8" resources), whichever occurs first.*
- *Under exceptional circumstances a member may submit an application for further assignments or allocations earlier than the expiration of this period.*
- *The APNIC Secretariat will monitor these exceptional requests carefully and publish comprehensive statistics on a regular basis. Without identifying any member organization, these statistics will record the numbers of requests and the outcome, the economy that the requests come from and clearly identify if any member has made more than one request under this provision.*

3.3 Conditions on recipient of the transfer

The recipient entity:

- *Must be a current APNIC account holder.*

- *Will be subject to current APNIC policies. In particular, in any subsequent APNIC IPv4 address allocation request, the recipient will be required to account for the efficient utilization of all IPv4 address space held, including all transferred resources.*
- *Prior to the exhaustion of APNIC's IPv4 space (that is, prior to the use of the "final /8" allocation measures) recipients of transfers will be required to justify their need for address space. After this time there is no requirement for any form of evaluation of requirements for eligibility.*
- *APNIC will maintain a public log of all transfers made under this policy.*

http://www.apnic.net/policy/transfer-policy

ARIN

8.3. Transfers to Specified Recipients

In addition to transfers under section 8.2, IPv4 number resources within the ARIN region may be released to ARIN by the authorized resource holder, in whole or in part, for transfer to another specified organizational recipient. Such transferred number resources may only be received under RSA by organizations that are within the ARIN region and can demonstrate the need for such resources, as a single aggregate, in the exact amount which they can justify under current ARIN policies.

https://www.arin.net/policy/nrpm.html#eight3

LACNIC

2.3.2.18- Transfer of IPv4 Blocks within the LACNIC Region

NOTE: This section will come into force when LACNIC or any of its NIRs becomes unable, for the first time, to cover an IPv4 block allocation or assignment because of lack of resources.

IPv4 block transfers shall be allowed between LIRs and/or End Users within the LACNIC region (hereinafter organizations) in accordance with the conditions set forth in this section.

2.3.2.18.1.- The minimum block size that may be transferred is a /24.

2.3.2.18.2.- In order for an organization to qualify for receiving a transfer, it must first go through the process of justifying its IPv4 resource needs before LACNIC. That is to say, the organization must justify before LACNIC the initial/additional allocation/assignment, as applicable, according to the policies in force.

2.3.2.18.3.- Upon receiving an IPv4 address block transfer request, LACNIC shall verify that the organization transferring the block is in fact the holder of said block according to LACNIC's records. The approved applicant and the organization transferring the resources must present before LACNIC a copy of the legal document supporting the transfer.

2.3.2.18.4.- LACNIC shall maintain a publicly accessible transfer log of all IPv4 address block transfers registered before LACNIC. Said log shall specify the date on which each transaction took place, the *organization from which the transfer originated, the receiving organization, and the block that was transferred.*

2.3.2.18.5.- The organization in which the transfer originated shall automatically be ineligible to receive IPv4 resource allocations and/or assignments from LACNIC for a period of one year as of the transaction date registered in the transfer log.

2.3.2.18.6.- A block that has previously been transferred may not subsequently be transferred again for a period of one year as of the transaction date registered in the transfer log. The same applies to its sub-blocks, which are blocks that group a subset of the IPv4 addresses contained in the block.

2.3.2.18.7.- Once the transfer is complete, LACNIC shall modify the information on the transferred resource in order to reflect the change of holder.

2.3.2.18.8.- The receiving organization must comply with all LACNIC policies in force.

2.3.2.18.9.- Blocks and their sub-blocks from allocations or assignments from LACNIC, being initial or additional, can not be transferred for a period of one year as of the allocation or assignment date.

2.3.2.18.10.- Transferred legacy resources will no longer be considered as such.

http://lacnic.net/documentos/politicas/manual-politicas-en.pdf

RIPE

5.5 Transfers of Allocations

Any LIR is allowed to re-allocate complete or partial blocks of IPv4 address space that were previously allocated to them by either the RIPE NCC or the IANA. Such address space must not contain any block that is assigned to an End User.

Address space may only be re-allocated to another LIR that is also a member of the RIPE NCC. The block that is to be re-allocated must not be smaller than the minimum allocation block size at the time of re-allocation. An LIR may only receive a transferred allocation after their need is evaluated and approved by the RIPE NCC, following the policies set for receiving further allocations within RIPE region (see the Section 5.3 Additional Allocations of this document).

Re-allocation must be reflected in the RIPE Database. This re-allocation may be on either a permanent or non-permanent basis.

LIRs that receive a re-allocation from another LIR cannot re-allocate complete or partial blocks of the same address space to another LIR within 24 months of receiving the re-allocation.

The RIPE NCC will record the change of allocation after the transfer. Please note that the LIR always remains responsible for the entire allocation it receives from the RIPE NCC until the transfer of address space to another LIR is completed or the address space is returned. The LIR must ensure that all policies are applied.

Re-allocated blocks will be signed to establish the current allocation owner.

Re-allocated blocks are no different from the allocations made directly by the RIPE NCC and so they must be used by the receiving LIR according to the policies described in this document.

http://www.ripe.net/ripe/docs/ripe-509

AFRINIC (Draft Policy)

2) The Proposal

2.1) Legacy members can transfer part or all of their IPv4 addresses to any company under the following criteria:

a) *The company to which the addresses are transferred may or may not enter into agreement with AfriNIC.*

b) *The legacy member may or may not inform AfriNIC about the transaction.*

c) *AfriNIC will accord the third party all relevant access to services and benefits normally available to legacy members.*

2.2) Paying AfriNIC members can transfer part or all of their IPv4 addresses to any company under the following criteria:

a) *The company to which the addresses are transferred must enter into agreement with AfriNIC.*

b) *The transfer and needs analysis cannot be based on any current policies. The only requirement for the transfer to happen should be the contract between the member and AfriNIC.*

c) *The relevant AfriNIC fees must apply to the third party.*

d) *AfriNIC will accord the third party all relevant access to services and benefits normally available to normal members.*

http://www.afrinic.net/docs/policies/AFPUB-2011-v4-001-draft-01.htm