Diffusion and Adoption of IPv6 in the ARIN Region

Hillary Elmore and L. Jean Camp and Brandon Stephens
helmore, ljcamp, bstephe @indiana.edu

1 Abstract

IIn the near term there will be no available, unallocated IPv4 addresses. From original estimates of IPv4 exhaustion in 2037,[12] the most widely-cited current estimates for ARIN IPv4 address depletion is now at 2013[13]. This deadline gives a particular importance to IPv6 adoption. The goals of this work were to identify valid measures of IPv6 diffusion and use classic diffusion models to bound the uncertainty in those measures. With these measures and simple models we can bound best case, current projection and reasonably optimistic cases for the adoption of the IPv6 protocol. For these ends, the work discusses previous analysis of IPv6 routes and ASN data from ARIN to quantify the current adoption rate. We conclude that there is no reasonable case for diffusion of IPv6 before IPv4 full allocation.

The second significant contribution, besides measurement and bounding uncertainty, that is provided in this paper is to what extent the now well established fundamental findings of the economics of computer security can apply to the diffusion of IPv6. The second significant but unanswered question is if the creation of a transferrable property interest in IPv4 addresses, informed by computer security economics, will hinder or galvanize IPv6 adoption. In order to address these questions the paper provides some non-trivial insights on IPv6 through presenting sketches of four scenarios: no action, IPv4 market creation, coordination government action and registrar- only management. As much as conclusions, this paper offers a set of questions that are critical to consider.

2 Introduction

In its first conception in 1977, it was believed that the address limits contained within IPv4 would never be an issue. Few but the most visionary [19] believed that Internet would expand the six orders of magnitude necessary to require the expansion of the IPv4 space. The transition to IPv6 is the price of once unimaginable success.

Now, not only do servers and microcomputers require IP addresses, desktops and laptops require connectivity. Increasingly mobile devices initially associated strictly with cellular networks require connectivity over IP. This increase in needed IPv4 addresses will exhaust the available address pools of the registrars within as little of four years. As the limit of unique IPv4 addresses (4,294,967,296) being approached1 it is critical to understand diffusion of IPv6. At its core IPv6 is a new addressing scheme developed in order to accommodate the continual expansion of the global network.[7]

IPv6 was initially introduced as far more than an expanded address space. Initially it was bundled with IPsec. Yet not only IPsec but also DNSSEC have been adopted over IPv4, to a greater degree than IPv6 itself. [1] Organizations adopted the elements introduced with IPv6 without moving to the new, larger address space.

Why has adoption of IPv6 been slow at best? We argue that this is due primarily to two well understood economic phenomena, without addressing in any manner the technical merits of the debate for or against IPv6. [2] These two phenomena are incentive misalignment and a lack of information. The lack of information is not only potential information asymmetry (e.g., a lemons market) but also lack of information with respect to risks and benefits for all parties. There is misalignment of incentives for many levels of adopters. At an

[1] As this paper makes clear, this does not imply widespread or arguably significant adoption at t his time.

[2] For example, the issues of the interaction of the address and the routing wrt IPv4 and IPv6 are clearly beyond the scope of this work.

organizational level, the more central a party is to the current IPv4 network the less incentive that party has to adopt IPv6. At a personal level, the more advanced a network engineer is in her knowledge of IPv4 the less she will encourage IPv6 adoption, as it could undermine her own expertise.

As the full allocation of the IPv4 space looms increasingly large, the need for an orderly transition to IPv6 is correspondingly critical. In this paper we use a simple diffusion analysis to provide a window into possible futures of IPv6 adoption. We find that there is, standing in 2008, no arguable diffusion path that will result in a seamless transition to IPv6. We argue that some of the reasons behind this are analogous to the lack of investment in other electronic networked risk (e.g., security) when investment often follows an eminantly foreseeable debacle, rather than being made to mitigate or prevent the event. We close by offering four scenarios: no action, transferrable rights over IPv4, coordinated governmental action, and registrar-ordered transistions.

3 Related Work In Diffusion

Technological diffusion has been studied in various disciplines for the past several decades, most notably in economics and sociology. [11] Bass developed the classic epidemic, information-based model in"'A New Product Growth for Model Consumer Durables". [2] This model assumes that, for all consumers except innovators, pressure to adopt a new technology increases as time and the number of other adopters increases. This model naturally lends itself to the development of an s- shaped diffusion curve; where diffusion of the technology first is spurred by innovators, but as the number of adopters increases their influence will diminish, as will the number of consumers who have not yet adopted the new technology. Thus, the rate of adoption slows as diffusion reaches its peak, completing the s-curve.

The model, grounded in epidemics, is most useful when studying the gradual impact of a new innovation. Knowledge of new technology takes longer to spread than, for example, knowledge of a world event because information on world events can be summarized, simplified, and broadcast from a common source. [10] New technologies also often have both hardware and software components. Hardware installation and factual descriptions of the protocol can be broadcast from a single source, and the information itself changes slowly. In practice, however, adopting IPv6 for a specific network and utilizing it in dynamic network conditions requires experience. The tacit knowledge that comes only from experience cannot be broadcast, but must be learned first-hand or through mentoring. Therefore, delay in adoption in the epidemic model can be attributed to the time it takes for the base of knowledgeable parties to reach a critical mass. Intuitively, simpler stand alone products are likely to diffuse more quickly than complex, integrated products.

The classic epidemic difussion model is as follows:

$$N(t+1) = N(t) + pN(t) + qN^2(t) \tag{1}$$

Here p is the innovator co-efficient, that is the rate at which early adoptors and innvoators adopt a particular technology. This diffusion curve has been applied (and proven) historically on televisions [15], telephones REF, e-commerce [25], various internet applications and a wide range of technologies [3]. Enhancements have enabled application of the basic model to a wide range of telecommunications product types. These have enhanced the flexibility of the model by building on its underlying structure, but not altered it. [8] Some of the variables to consider are either currently unknown in the case of IPv6 (price, advertisement) or suggest the possibility of future research (e.g., type of user, market size).

Though the epidemic (or population) model of diffusion is the widely use, it may not be the most applicable in all situations and for all purposes. The probit model of diffusion examines not cumulative diffusion, but the diffusion to individual firms. [10] To do this, the probit model assumes that every firm sets a threshold of profitability. When profitability of the new technology is below the threshold, the firm will not adopt the innovation. When the profitability rises above this internally defined threshold, the firm will choose to adopt the new technology. The S-curve in this model is determined by the change in profitability of a particular technology and firms' changes in their threshold for adoption as more information about the technology becomes available. Therefore, the probit and epidemic models differ in that the epidemic model assumes that adoption of a new technology will occur when a firm becomes aware of it, while the probit model explains, to some degree, the apparent hesitation of firms to adopt new technologies even after they are aware of the technology. The probit model and the S-curve are not mutually exclusive. This is

particularly true when profitability (benefit) is a function of the number of previous adopters (e.g. network effects) as is the case with IPv6.

Though the literature on technological diffusion makes a strong case for the S-curve model, like all models it has flaws. In particular, though diffusion can be studied retrospectively, using the S-curve model to predict and prescribe new technologies is problematic, particularly when applied within individual firms. [6] Improvement in a technology is difficult to predict, and is sensitive to external factors. Simply forecasting that accepted technology is approaching its natural improvement limit can have a direct downward effect on the technology's growth trajectory. [6] This effect both causes and is furthered by a subsequent decrease in engineering resources devoted to the displaced technology, as well as an increase in resources toward the innovative technology. In this case, the exhaustion of IPv4 has been announced several times, but new technologies have been found.

Much of the diffusion literature is built upon two basic assumptions: first, that new technology (once it is released) and the old technology do not change during the diffusion process, and second, that the new technology is better than the old technology. This first assumption is often proven wrong in the real world, since the quality of the new product does, in fact, increase during diffusion, and that this improvement causes the equilibrium adoption point to rise continuously. [5] Hall argues that it cannot be assumed, either, that the old technology does not change during the diffusion of the new technology, as old technologies can experience a 'last gasp' improvement in an effort to remain dominant. [11] The old technology, faced with competitive pressures or simply with its own limits, also increases in quality. This has certainly been observed with IPv4, for example, with DHCP increasing flexibility of internal address allocation. Changes in either or both of the technologies delay adoption of the new technology because changes increase the uncertainty of the benefits of the new technology and can increase the risks, especially if the old technology is able to incorporate aspects of the new. Again, this applies to the case at hand. One of the initial perceived benefits of IPv6 was its linkage to IPsec and the fundamental benefit was the ability to provide more addresses. IPsec has been unbundled from IPv6. NAT and DHCP have enabled effective sharing of IPv4 addresses. Thus the benefits of switching have been reduced.

4 Data Experiment Setup

The major question that this analysis seeks to answer is: given current adoption rates, might IPv6 have significant domestic market penetration before the exhaustion of the ARIN pool of unassigned addresses? The answer, not surprisingly, is no. The second question on the size of the gap between IPv4 allocation and IPv6 adoptions.

This analysis was originally focused on attempting to see which factors were important in determining why a firm was or was not adopting the IPv6 protocol. This was to be done by looking at the ARIN Project data and filtering out firms in different sectors of the market. After this was performed, criteria would be set up to decide why a firm was adopting as compared to all other firms in that sector. Unfortunately after analyzing the data, we realized that there was not sufficient multi-sector adoption to prove any worthy prediction of the market factors. Domestic adopters consist almost entirely of network service providers and the US government. Content providers, e-commerce sites, and hosting companies have not adopted. Probit analysis for the ARIN region therefore seems premature.

We have therefore adopted a macro S-curve approach to IPv6 adoption.

The first critical choice was the determination of the variable to use for IPv6 adoption. Traffic data are likely to be unrepresentative, and may require human subjects approval. While Indiana University is home to the Internet2 NOC, the data that are available may be quite unrepresentative of the larger network. Obtaining reliable representative traffic data is a significant problem for many types of network research.

The second choice of data source was an evaluation of open IPv6 routes. This preliminary analysis was deemed untrustworthy after some additional consideration for three reasons. First, an IPv4 route is not equivalent to an IPv6 route. Second, the diffusion model assumes that there is a static overall pool of potential eventual diffusion. That is, 100% remains 100%. Even with removing duplicate routes (i.e., those with dissimilar endpoints) the IPv4 data were considerably more fluid than the IPv6 data. This injected even more uncertainty into an already uncertain issues.

There are two serious limitations to the use of routes to compare diffusion of IPv4 and IPv6. First, the

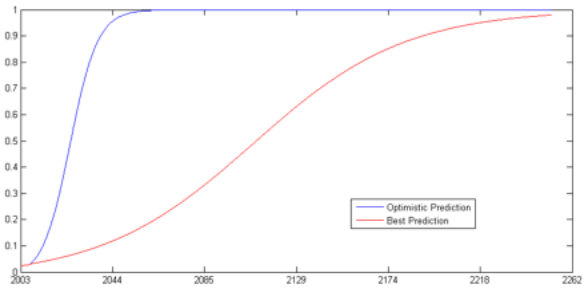

Figure 1. Best Fit with ASN Data

same number of routes does not reflect the same level of adoption, e.g. it is to compare apples to oranges. The second issues is that advertized routes, are by definition, public. One reason to adopt IPv6 is because it may be able to provide a better ability to determine what devices are 'inside' the trusted network. Thus an organization may choose to use IPv6 internally without advertising its routes.

One reason to adopt IPv6 internally is to address the issue of network boundaries. By providing each device with its own combination of machine address and unique IPv6 address, it is possible (in theory) to observe the addition and removal of devices with more certainty. Porous network controls from potentially hostile laptops, mobile devices, and even photo frames are an increasing problem for corporate security. Firewalls are inadequate for confirmation of a trusted insider versus a trusted outsider.

Thus, we concluded that a second source of data was needed. The follwoing analysis uses the same equasions, and in the best case the same data truncation, as we initially used with route data. [9]

The data we selected was Autonomous Network System numbers. We determined that one ANS with IPv6 was far more like one ANS with IPv4. The results indicated that even in the next possible scenario, within the possible range of error and with data selection most favorable to IPv6 adoption, there will still be a multiple year transition period. In the worst case, it is feasible that IPv4 and IPv6 co-exist on the network for many decades.

The argument for this analysis is that one ANS for IPv4 is equivalent to an ANS for IPv6. The counter-argument is that obtaining an assignment does not imply actaul use.

5 Data Analysis

Predicting the future growth of the IPv6 adoption rate by the current data using best-fit results in the S-Curve shown in Figure 1. The data corresponding to Figure 1 shows that at the current rate of adoption, it will take approximately 15 years for a 50% adoption of the IPv6 protocol. As for an 90% implementation of IPv6, it will not occur until 2044 in the best case. Unfortunately with fewer and fewer IPv4 address remaining available, IPv6 will need to be adopted far before these dates. Notice that what is shown is an envelope of possible adoption. The left hand curve[3] shows the results with the follower coefficient (i.e., q) increased by one standard deviation. The right hand curve shows the results with the follower coefficient decreased by one standard deviation. Thus, the fifteen and thirty year windows are the most positive possible interpretation. It as as likely that IPv4 and IPv6 will co-exist through the lifetime of the network, as shown in the right hand curve.

This analysis does not take into account demand push (i.e. exhaustion) but only supply pull. Resource exhaustion is a significant component in IPv6 diffusion globally. According to ICANN2, Mexico's DNS distributor will stop allocating addresses January 2011, and many other countries are following close behind in similar policies. Therefore, systems will need to be in place much more quickly then the current domestic adoption rate can manage.

One flaw in the use of standard diffusion curves for the diffusion of IPv6 is that the exhaustion of IPv4 may cause a dramatic increase in IPv6. (Simultaneously, it may cause an increased investment in technologies

[3]The left hand curve is in blue, if this is seen in color.

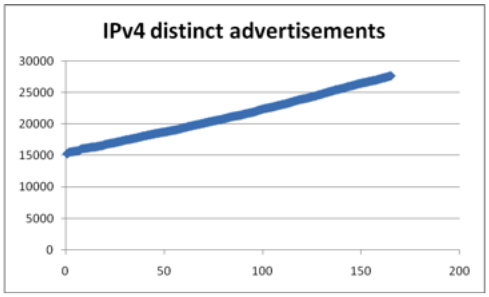

Figure 2. IPv4 ASNs over Two Years

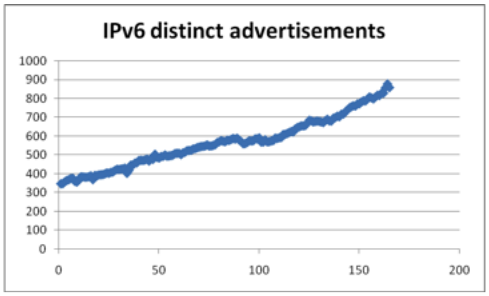

Figure 3. IPv6 ASNs over Two Years

to leverage IPv4, e.g. next generation NAT). Predicting such a thing leaves the realm of modeling and moves into the realm of sheer assertion. This implies adding almost arbitrary data, and once data become arbitrary we have left the realm of analysis for simple assertion. So how to combine a data-inspired approach with the unknowable implications of exhaustion? We simply have chosen to include the uncertainty in the results, rather than trying to provide a pretense of knowing the uncertain.

There are two reasonable ways in which to adjust the route-based findings. The first of these is adding exogenous data for DoD (or similar) adoption points as forcing functions. Adding these to the pre-existing data yielded nearly impossible curve fits within the standard equation. Indeed, in multiple attempts to fit the curves we had negative innovator coefficients. (We were able to force the data using route numbers, resulting in the best case number referenced above, of 80% diffusion in 8 years.) So the second change was to truncate the previously used data so three months.

When is it reasonable to consider IPv6 diffusion as initiating? For example, the first implementation that we would recognize as a fax was sent in England in 1843, and telephotography was reliably demonstrated in England in 1902. The first transcontinental fax was sent in 1955. Yet a diffusion study of fax would not reasonable begin until the devices came into mass production for sale outside IBM, ten years later. There is a question as to the study of IPv6 diffusion should reliably began.

That IPv4 is well-established is clear. There is an argument to be made that IPv6 did not truly begin to diffuse until after the 6bone termination project. Before that time, much of the adoption is a result of that project. What is apparently termination of the 6bone project appears clearly in Figure 4 as a significant and apparently sudden drop to beneath 0.02.

At the end of 6bone project, all adoption was evidence of diffusion. Thus a reasonable analysis is to truncate the data to the last five months of 2007 and January. The results are shown in Figure 5.

In Figure 5 the data are again shown as an envelope encasing the possible range of adoption points. Again the left-hand side has the follower coefficient (i.e., q) increased by one standard deviation. The right hand curve shows the results with the follower coefficient decreased by one standard deviation. Now, instead of showing a possible range of 2044 to 2240 for adoption, the range becomes 2014 to 2080, between six and seventy years.

Even though the adoption timeline of six years is best case, this may be inadequate in terms of responding

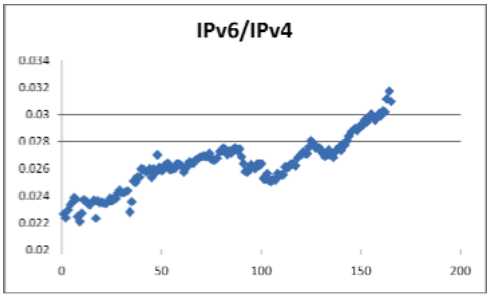

Figure 4. IPv4 v IPv6 Ratio of ANSs

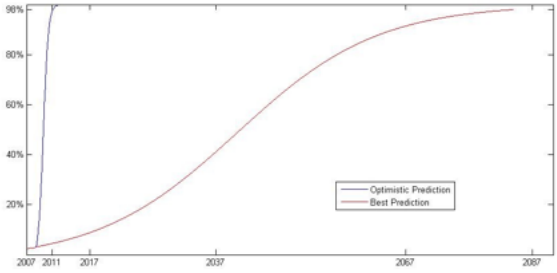

Figure 5. Diffusion using ANSs and Truncated Data

to IPv4 exhaustion. Recall the projected date of IANA's unallocated address pool exhaustion is November 28th, 2010, along with RIR's projection as being December 5th, 2011.[13] At this point, two of the major sources of IPv4 address allocations will no longer have a pool of IPv4 addresses to assign.

It is worth noting that even without the efforts to develop an absolute best case of adoption, the adoption co-effecients are extremely high, and do not reflect the adoption of Internet use, television use, radio use or any other communications technology. These coefficients are arguable in that they can be compared to adoption rates of applications, e.g. fetch or http.

The difference between the two figures illustrate that the uncertainty in IPv6 diffusion is great. The results are sensitive to the initial conditions. That initial conditions have tremendous influence to the outcome is not surprising to anyone who has ever run a model. We argue that going back further than or as far as shown in Figure 1 would be difficult to justify. A strong argument can be made in qualitative terms that the closure of the 6bone project is the beginning of IPv6 diffusion.

Recall that the decision to use routes was rejected, and we determined that we should instead use ASNs. Note that the results from the IPv4 route data were not wildly dissimilar from the results from the ASN data. In the route data the best fit resulted in 80% adoption in 22 years,and the most optimistic that can be extrapolated with current route data is 80% adoption in 8 years. [9]

We have not endeavored to use our models to create false certainty, nor have we pretended omniscience. However, we have used the best available data to constrain that uncertainty, and offer a straight-forward method that anyone can duplicate. In any case, the results show that there will be a nontrivial, multi-year window between exhaustion of the unassigned IPv4 address pool and the widespread adoption of IPv6. The final illustration is the most powerful. Given a solid argument about the actual adoption of IPv6, we have illustrated that it is possible that it will be decades while IPv4 and IPv6 coexist. The following sections argue about the sources of this delay, and the possible implications as we move further from the data.

6 Related Work In Economics Of Information Security

Many of the most significant issues in the study of the economics of information security also apply to the adoption of IPv6. In order to understand why the creation and specification of IPv6 has yet to result in its widespread diffusion, we began by taking an economic approach to analyze possible sources of the current adoption rate. We focus on the misaligned incentive structures that are faced by IPv6 adopters. In doing this we believe that we can determine a partial cause to the slow adoption of the IPv6 protocol. For example, for information security to be effectively adopted, incentives must be properly aligned and must not allow hidden actions. [1] Network externalities are also applicable to the (lack of) adoption of IPv6. An increase in the size of the network increases the value of the network to each user. As a result, a small network of deployment results in small total benefit. With the small network, the cost of adopting a new technology is greater than the benefits gained by the added protections of the network. In marketing terms, the technology may never reach critical mass. Moores Law applies to security as well as networks.[24]

Patching behavior literature is generally based on the fact that not everyone who can apply a patch does so. Cavusoglu et al. note that there are four main reasons individuals and firms fail to apply patches. [4] First, too many vulnerabilities exist to individually apply patches. Secondly, patches will not be applied until they are trusted, and they cannot be trusted until they have been tested on a system. Third, there is no standardization in the distribution of patches, and finally, patches require testing after installation.

These difficulties in patching apply, in many ways, to the difficulties in IPv6 adoption. The most obvious parallels are to the testing, trust, and installation aspects of patches. Though most routers are now sold with IPv6 capabilities, the move from IPv4 to IPv6 could still cause disruptions in service, as well as temporary increases in security vulnerabilities. [23]. Regardless of the capabilities with which a device is sold, IPv6 must be enabled on devices for them to work. A product advertised with a full IPv6 stack may not have a fully functional stack. Because of the complexity of IPv6, two implementations may be compliant but suffer subtle failures in interoperability. The increases in switching costs to IPv6 coupled with the bootstrapping effect provide a plausible basis for the current slow adoption of IPv6 in the United States. [23]

Adoption of new protocols in the technology world can be very costly, not only in monetary terms, but also in time spent understanding and deploying the technology. [23] In"Could IPv6 Improve Network Security?", Rowe presents estimates for the costs of implementation and the benefits gained from the implementation of IPv6. Rowe begins by estimating the incremental cost of labor and training required for the IPv6 conversion at approximately 25 billion dollars. This amount would be spent over an estimated 25-year implementation period, which seems like a very large cost for the system, but it is negligible when compared to the hardware and software costs associated with the conversion, less than 1% estimated. [23] . This large discrepancy in the cost of adoption versus benefits places a large burden on initial adopters of IPv6, one that most companies cannot bear in terms of maintaining comity with stockholders. Ironically, IPv6 may increase near-term security vulnerabilities because of the relative immaturity of the software. [22] In addition to the inherent problems of a younger code base, lack of employee experience may increase misconfiguration. Misconfiguration is an extremely common event. [20] A major cause of real world vulnerabilities is the unintended interactions of software components. Each component may be secure independently but create vulnerabilities through their interactions. [18] All of these costs weigh heavily on the shoulders of early adopters. Unfortunately, as stated by Jaffe, 'the initial benefits obtained by early adopters might fall significantly below the costs of adoption.' [14] This can be a large negative incentive for early adopters as they tend to incur most of the cost and tend to have to wait for long term gains for a return on their investment.

Market interventions to offset the costs incurred by early adopters promote adoption. In subsection 5.3 we modeled three adoption paths with the same type of forcing function. This can be done several ways. The first that we will discuss are subsidies, or a grant provided by the government. Subsidization of technology-based adoption is not uncommon in the global economy, and even though it is not currently being done domestically for IPv6, it is being done globally. Take, for instance, South Korea, China, Japan, and the European Union

Another form of cost displacement is a fine. Fines allow the adopter to 'receive a conceptual benefit' by complying and not having to pay the fine. [21] Fines are commonly used as a negative incentive. This in turn causes them to not produce the full expected effect as people respond better to positive incentives than to negative incentives. Content-based adoption incentives have been implemented in a few markets.

However, these have proven ineffective.

By providing incentives like those mentioned above, the United States may be able to induce more adopters to enter the market. This in turn will drive prices down as there will be a larger supply of products including software, hardware and support for IPv6. As the prices go down, the demand will continue to increase until it hits near market saturation, either by simple replacement of old hardware, or by actively adopting. This will of course take time, but with positive incentives to adopt, typically supported through strong governmental policy, the rate of adoption can cause the adoption timeline to become relatively short.

7 Implications

In this section we discuss four high level views of possible futures. Our high level conclusions are not encouraging, but defensible. Taking no action is not tolerable, and will become increasingly intolderable if entrepreneurs are denied network access. The exclusion of innovators and competitiors from the network will result in governmental action. Providing IPv4 rights can create transparency, yet designing a property right and transfer mechanisms that motivate IPv4 adoption rather than unintended strategic behaviors requires more than a single stroke of brilliance. Coordinated governmental action to avoid diffculties in transition is optimal, yet somewhat less likely. We conclude that the registrars themselves must take action to manage the transistion, and offer a single proposal more as a straw man than as marching orders.

7.1 No Action

On a global scale, the benefits of IPv6 adoption have the potential to outweigh the costs for developing and late-adopting nations in the near term. It is also appears to be the case that AfriNIC will be the registry with the last remaining IPv4 addresses available for allocation. At that point, AfriNIc may be able to choose to to leverage the IPv4 addresses by requiring some commitment to African Internet development. (Or as discussed below, sell these on the proposed IPv4 market.) Developing nations stand to see significant benefits from developing IPv6 infrastructure at this point.

Given the current expenditures on IPv4 in the United States and the investment cost necessary to switch from IPv4 to IPv6, this may not be the best option for the U.S. and other developed countries with existing IPv4 infrastructure.

Though IPv6 purported to address many of the security flaws of IPv4, the fact that IPsec has subsequently been applied over IPv4 limits the benefit of this aspect of IPv6 adoption. In addition, as Rowe suggests, the transition to IPv6 will inevitably result in unforeseeable new security vulnerabilities. [23] The marginal benefits of IPv6 over IPv4especially since many of the security enhancements of IPv6 have been implemented over IPv4and the high switching cost, there is an argument that it will remain beneficial for U.S. companies to continue operating over IPv4.

European authorities, even less than American regulatory authorities, are unlikely to tolerate a situation where incumbents are able to prevent interconnection through their own failure to adopt new technologies. Forced adoption would be a likely long term but difficult and contentious regulatory battle. The level of deployment in Europe was termed"imperceptable" in the final 2004 report of the European IPv6 Task Force.

The U.S. may choose to effectively remain alone as the world converts, as with the case of the English to metric conversion.

7.2 IPv4 Market Rights

The unanswered question is if there is a price point for IPv4 addresses which will drive IPv6 adoption. In order to enable an orderly transition, there is initial research in enabling the transfer of IPv4 addresses through the registrars.

The creation of a functional market in IPv4 addresses requires, as a minimum:

1. a mechanism for clearing the market at the appropriate price,
2. an ability to ensure exclusive use of an IPv4 block once allocated,
3. a bundle of rights that will be transferred, and

4. mechanisms for dispute resolution.

The creation of a functioning market, without opportunities for arbitrage, collusion, or the creation of hazardous outcomes has proven problematic in the related realm of spectrum auction design. IPv4 addresses are similar to spectrum in that these are required for companies to either go into or remain in business. As a cellular company either has a spectrum license or no business model, an Internet network services company either has access to a routable reachable address or no business. [17] The existence of a market in and of itself does not solve issues of competition, which are, 'preventing collusive, predatory, and entry deterring behavior'. The last decades of spectrum license auction design have shown these to be difficult design goals.

What of the dispute resolution mechanisms? What bundle of rights is an IPv4 address? Is there a requirement that the addresses be routed or routable? If so, what party is responsible for ensuring that the addresses will be routed or routable? The choice of routing an IPv4 address may be made by a party which is competing for an IPv4 block. Certainly the registrars cannot assure that a block is routable, except by providing a large enough block. If there is a dispute between two parties, so that both are actively publishing the address, much of the cost will be borne by legally unaffiliated but physically networked parties. Use of IPv4 addresses for internal routing in a time of scarcity may not be socially optimal, but may prove quite affordable. How can this predictable response to a conflict in ownership be avoided?

The ability to prevent illegitimate use, that is the right of exclusion, is a fundamental property right. Can the registrars enable or support the rights necessary for reliable transfers of IPv4 addresses between parties? Currently the contractual authority of the registrars is untested as there has never been a case of a revoked IPv4 address.

There is far from global consensus on the legitimacy of ICANN, and, as with the domain name registrars, any act by the registrar to recall or revoke an IPv4 address will certainly be tested in the courts. [16] While as a whole, it is not in the interest of the registrars and network operators to appeal to the courts, under an IPv4 market regime it is likely that such an appeal will be in the interest of one party at least one. Unlike the case of DNS, there is no root for the registrars to leverage to prevent messy legal conflicts from generating technical disorder.

The largest holder, without question, of IPv4 space is the US Government, particularly the Department of Defense. With the least possible stimate of the base budget for the DoD in 2007 being half a trillion dollars, it is unlikely that the price of a IPv4 address will be large enough to inspire entrepreneurial action within the military ranks. However, it is possible that a market cound encourage a small number of well-networked individuals to leverage the arbitrage opportunity through regulatory manipulation, thereby releasing some of those addresses on the market.

Any market that is designed will create opportunities for regulatory or market arbitrage. At its foundation, market arbitrage is the standard practice of selling what you may not own in order to purchase the good at a profit. Regulatory arbitrage refers to crossing regulatory boundaries for profit. For example, an initial incentive for voice over IP was the avoidance of international telephony rates, many of which were set by government- owned PTTs. Regulatory arbitrage with an IPv4 market may include registrar shopping (seeking an optimal transfer regime) or jurisdictional shopping for dispute resolution.

Designing an optimal market which spurs the adoption of IPv6, rather than enabling undesirable market behaviors, is at best an extraordinarily difficult task. In the absence of gaurantees of functionality (i.e. address will be routed), without the capacity for exclusion, and with unspecified dispute resolution, the problem is exacerbated.

Finally there is the question of what to do should the IPv4 market be a perfect success. The primary function of a market is to manage scarcity. If the market is successful, then the scarcity of IPv4 addresses will be managed and the incentive to adopt IPv6 dissipate. To the extent that a market in IPv4 enables long term management of that scarcity (with or without addressing barriers to entry, and predatory pricing) IPv6 may never come to fruition.

7.3 Coordinated Governmental Action

Note that one conclusion of this study is that in the ARIN region, IPv6 is not being adopted at a rate comparable to other countries, according to the claims of the other nations.

The current high switching costs and low perceived benefits of switching are discouraging the adoption of IPv6 among service providers and early adopters. Governmental supportin the form of subsidies, training,

demand pull, information provision and policy changes along with possible tax cutscould provide incentives for early adopters to switch to IPv6. However, as seen from the data analysis, even this will not adequately increase adoption rates and enable the US to implement IPv6 in a timely manner.[4] As the estimate for IPv4 exhaustion is uncertain, it is difficult to determine the exact time at which no IPv4 pool will exist to support allocation of more addresses. The one certain conclusion of that work is that this exhaustion time will occur before IPv6 adoption.

Should investment be made to force acceptance of IPv6, which is a tightly bundled product? Or should US and European governments immediately invest in research on potential IPv4 expansion alternatives, e.g. daughter of NAT? Without action on adoption or unbundled alternatives Europe and the US risk having chronically insecure, unnecessarily expensive, second generation IP networks. With overly aggressive action, the US and Europe could force premature adoption causing a window of greater disruption and vulnerability.

One obvious possibility, in particular for the US government, is to announce and then provide for the release of IPv4 addresses to the registrars. The US DoD is in the unique position of having adequate IPv4 dark space that is possible for the DoD to negotiate in reasonable faith with the RIRs. Other governments have the ability to tax, regulate or offer monetary incentives. Only the US DoD has the ability to expand the available IPv4 space by freeing any unused darknet.

The efficacy of lobbist-driven release of IPv4 adresses on a market (see the previous section) versus releases of IPv4 to the registrars in the national interest is primarily an ideological rather than an empirical question. The answers to that particular question in both dimensions are beyond the scope of this work.

7.4 Registrar - Ordered Transition

If it is the case that IPv6 is the protocol that must be adopted to ensure an open Internet, then the RIR community has choices to make about how that might happen.

An obvious choice is to alter allocation of IPv4 blocks so that each allocation is the minimal size to fit in the routing architecture. That is, allocate no party more than a /20. This would require refusing additional allocations to those parties which already have received IPv4 allocations.

An alternative choice is to only allocate IPv4 addresses to organizations without existing allocations. If IPv4 has, in the long term, a role as a translator to what will become the legacy network, the argument for additional allocation is weak. If organizations which already have IPv4 blocks which can be routed are assigned only IPv6 addresses, this implies that the most rapidly expanding entities on the network will have the greatest incentive to move to IPv6.

Making these choices is made more complex by the fact that the RIR communities consist exactly of those organizations which already have IPv4 blocks. Thus the RIR will effectively be asking its membership to deny itself access to potentially valuable address space to ensure that others have this address space. Alternatively, the RIR will be asking its members to maintain the current governance mechanisms by insuring that there is neither a requirement for market or additional regulation.

The RIR community may still make this choice although the window for effectively making such a decision is closing. The choice will burden members by requiring IPv4 translation in the near term and a move to internal IPv6 space as opposed to obtaining larger allocations.

8 Closing

IPv6 adoption either will not happen in a timely manner, and should be encouraged by some combination of regulatory force and some class of incentive. Simultaneously, it is necessary to plan for a significant period of overlap, where both IPv4 and IPv6 coexist.

Having implemented the cursory examination of the problem required to obtain the estimates for IPv6 diffusion, it is clear to the authors that there is a commons problem in IPv4 exhaustion. It is in the interest of all parties to maintain address allocation as a technical rather than a political process. The currently involved parties make technical arguments for more space, and the individuals determining standards do so

[4] In this case, a timely manner would mean in time to avoid a period of IPv4 exhaustion before IPv6 adoption.

from an informed engineering basis. [5] Essentially what is required is for the RIRs to solve their coordination problem or to prepare for inevitable regulation, regulatory body instantiation and contention over IPv4 use, rights and even 'ownership'.

As history illustrates, no incentive is as great as the market compulsion of resource exhaustion. While nothing focuses the mind as well as the prospect of a hanging in the morning, this applies only to the hung. As long as the prospect of IPv4 allocation provides opportunities for prohibition of entry, competitive advantages in provision of addresses, and limits on competition those organizations at the core of the network have arguably perverse incentives with respect to adoption of IPv6.

Some combination of government encouragement of IPv6 development in the form of subsidies, adoption mandates, and bundling of technologies could help increase adoption rates. Yet previous economics of security work has suggested and time has illustrated that information provision, subsidies, training and demand pull have proven inadequate for eliminating even those vulnerabilities based on basic logical errors. Arguably, because IPv6 technology is present in most routers sold today, simply making it availableas in the bundled technology strategyis inadequate to truly hasten its adoption. Subsidies and mandates together might prove sufficient incentives for firms to switch to IPv6.

We have begun work on evaluating whether the domestic timeframe is concurrent with international adoption timeframes, or how far ahead of or behind global adoption the domestic market is.

References

[1] R. Anderson and T. Moore. The economics of information security. *Science*, pages 610–613, 2006.

[2] F. M. Bass. A new product growth for model consumer durables. *Management Science*, pages 215–227, 1996.

[3] F. M. Bass, T. V. Krishnan, and D. C. Jain. Why the bass model fits without decision variables. *Marketing Science*, pages 203–223, 1994.

[4] H. Cavusoglu and H. C. and Jun Zhang. Economics of security patch management. In *workshop on the Economics of Information Security*, Cambridge, UK, June 2006.

[5] G. C. Chow. Technological change and the demand for computers. *American Economic Review*, pages 1117–1130, 1967.

[6] C. M. Christensen. Exploring the limits of the technology s-curve. part i: Component technologies. *Production and Operations Management*, 1(4):334–357, 1992.

[7] S. Deering and R. Hinden. Rfc 2460 - internet protocol, version 6 (ipv6) specification. http://www.faqs.org/rfcs/rfc2460.html/, 1998.

[8] C. J. Easingwood and S. O. Lunn. Diffusion paths in a high-tech environment: clusters and commonalities. *RD Management*, page 6980, 2002.

[9] H. Elmore, L. J. Camp, and B. Stephens. Ipv6 in our lifetime? http://www.iepg.org/november2003/, 2008.

[10] P. A. Geroski. Models of technology diffusion. *Research Policy*, pages 603–625, 2000.

[11] B. H. Hall. Innovation and diffusion. In *Handbook on Innovation*. Oxford U Press, 2004.

[12] G. Huston. Ipv4 address lifetime expectancy revisited. http://www.iepg.org/november2003/, 2003.

[13] G. Huston. Projected iana unallocated address pool exhaustion. site: http://www.potaroo.net/tools/ipv4/, 2008.

[5] This is in no way a claim that there are no conflicts of interest within Internet standards bodies. That is beyond the scope of this argument.

[14] A. B. Jaffe, R. Newell, and R. Stavins. A tale of two market failures: Technology and environmental policy. *Ecological Economics*, 54:164–174, 2005.

[15] M. Karshenas and P. Stoneman. A flexible model of technological diffusion incorporating economic factors with an application to the spread of colour television ownership in the uk. *Journal of Forecasting*, pages 577–600, 1992.

[16] J. Kesan and R. Shah. Fool us once shame on you - fool us twice shame on us: What we can learn from the privatizations of the internet backbone network and the domain name system. *Washington University Law Quarterly*, pages 89–114, 2001.

[17] P. Klemperer. What really matters in auction design. *The Journal of Economic Perspectives,*, 6:169–189, 2002.

[18] C. E. Landwehr, A. R. Bull, J. P. McDermott, and W. S. Choi. A taxonomy of computer program security flaws. *ACM Comput. Surv.*, 26(3):211–254, 1994.

[19] J. Licklider and R. W. Taylor. The computer as a communication device. *Science and Technology*, pages 21–31, 1968.

[20] R. Mahajan, D. Wetherall, and T. Anderson. Understanding bgp misconfiguration. In *SIGCOMM '02: Proc. of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 3–16, New York, NY, USA, 2002. ACM.

[21] A. Ozment. Bootstrapping the adoption of internet security protocols. In *Fifth Workshop on the Economics of Information Security*, Cambridge, UK, June 2006.

[22] A. Ozment. Milk or wine: does software security improve with age? In *USENIX-SS'06: Proceedings of the 15th conference on USENIX Security Symposium*, pages 7–7, Berkeley, CA, USA, 2006. USENIX Association.

[23] B. Rowe and M. Gallaher. Could ipv6 improve network security? if so, at what cost? *I/S A Journal of Law and Policy for the Information Society*, 2006.

[24] S. Smith. Magic boxes and boots: Security in hardware. *Computer*, 37(10):106–109, 2004.

[25] J. Ure. Modelling critical mass for e-commerce: the case of hong kong. *Electronic Commerce Research*, 2(1-2):87–111, 2002.