Economic factors in the allocation of IP addresses

Dr. Milton Mueller, Syracuse University School of Information Studies, USA and XS4All Professor, Technische Universiteit Delft, Netherlands. With the assistance of Youngseek Kim, Ph.D. candidate, Syracuse University School of Information Studies

SUMMARY

This study reviews the economic rationale for IP allocation policies and argues that conservation is still a valid concern for IPv6 due to the large basic unit of IPv6 allocation. According to the study, current RIR policies provide few incentives to reclaim unused or under-utilized resources. The study concludes that it is possible to have a much more liberal policy for the initial stage of IPv6 allocation, provided it is associated with efficient reclamation policies and incentives, which incentives should be economic.

After analysis of the IPv6 fee structure charged by the RIRs and brief discussion of the routing scalability issue of IPv6, the paper proposes a Transferable Address Block Lease (TABL) model as a market-oriented allocation mechanism. The TABL model employs two economic techniques: (1) fees related to the size of the IPv6 address blocks and (2) transferability. TABLs would be an additional option, not a replacement for needs-based allocations, and confined to the mid range of IP address blocks (/48 - /32). The paper discusses auctions as a method of making initial allocations and concludes that they would not be advisable.

The chief advantage of TABL is that it makes addresses available to anyone who wants them in the quantity they need, thus increases end users' control over their own network destiny and reduce their switching costs to change providers. However, the market mechanism embedded in this model doesn't directly address the latecomers' fairness concerns.

The concern that the absence of needs assessment might lead to wasteful allocation or inefficient usage of IPv6 resources is answered as follows: charging fees related to the size of the block would provide the right conservation incentives, and the number of TABL blocks any one organization/entity could lease would be limited. Moreover, the combination of fees and transferability would facilitate shifting resources to their most highly valued uses. The concern that TABL could result in larger routing tables is avoided by restricting disaggregation of the TABL blocks. Practical issues in implementing this model include setting the initial allocation fees and the rate at which fees would rise based on size. However, the study concludes that these problems could be solved through an incremental approach and it suggests to use the RIR fee structure as a reference point and choose a fee structure that makes a TABL slightly more expensive but still attractive.

The study suggests that ITU could play a constructive role in a possible implementation of TABL. With its experience and expertise in international legal and regulatory frameworks, ITU could develop a model framework for the TABL contract. TABL could be implemented by ITU on a global basis or it could be implemented in countries who wish to adopt this model for their national allocations under the CIR model.

Final version (2.2), 9/17/2009

Contents

Economic factors in the allocation of IP addresses

This report analyzes the potential for market forces to supplement or substitute for certain RIR functions. The purpose of this exercise is not necessarily to advocate market-oriented policies, but simply to gain a better understanding of what options exist and what implications they might have for address management. The discussion is divided into four sections. The first section analyzes the techno-economic characteristics of IP addresses and their interaction with routing in abstract terms. The second section focuses on the IPv6 address space. It analyzes the IETF recommendations, RIR policies and fee structures pertaining to IPv6, and identifies key questions about scarcity, routing and allocation policy. The third section explores how this newly-emerging IPv6 institutional regime might be changed to employ economic factors (pricing and markets) in the allocation of address space. The last section assesses the possible advantages and disadvantages of the proposal in relation to the status quo, and the possible role the ITU might play in facilitating or implementing such changes.

1. Techno-economic features of IP addressing/routing

This section links the analysis of IP address policy to the established vocabulary and concepts of institutional economics. Addressing and routing are usually discussed in technical terms, yet embedded in this highly technical discourse are a number of critical economic concepts, such as scarcity, externalities, tragedy of the commons, and the distribution of costs. Aside from the pioneering work of Geoff Huston, an Internet engineer who makes occasional forays into the engineering economics of Internet service,[1] there is very little research literature that attempts to bridge economic, institutional and technical discussions of IP addressing and routing. IP addressing/routing studies are just beginning to enter the phase that studies of Internet and information security began about a decade ago, in which the dependency of technical operations on economic incentives is acknowledged and the influence of economically-motivated behavior on the design of technical systems is integrated into research and policy analysis. (Anderson, 2000)

1.1. General observations

If we look at the IP address space in isolation, the problem we are confronted with is not very complex or unusual. The IP address space creates a virtual resource of finite dimensions. It is similar to a block of radio frequencies dedicated to a specific service by technical standards. IP address space size is fixed by the technical standards defining the Internet protocol. IP addresses are scarce in the strict sense that economic theory defines scarcity: it is not possible for all of us to have all of the addresses we would like at zero cost. As a virtual resource, they are not "consumed" or used up when put into production; rather, they are occupied just as radio spectrum or land parcels are occupied. When the occupation of one party ends, the resource could be available for others to occupy. In Internet discourse, occupying an address is called an "assignment;" the occupation of a larger contiguous block of addresses by an intermediary for assignment to others is called an "allocation."

This relatively simple picture of a resource space breaks down when routing is brought into the picture. Routing is the process that guides the movement of Internet

[1] (Huston 2001, 2005, 2009)

protocol packets from their origin to their destination. The policy problem we are faced with is not, in fact, one of efficiently allocating IP addresses *per se*. The overarching problem is actually the efficiency and scalability of Internet routing. The possession of IP addresses is significant and valuable only insofar as they can be used to route packets on the public Internet; likewise, routing needs globally unique IP addresses in order to define routes that can guide packets to their destination. The resource space we are confronted is complicated by the interdependence of routing and addressing, which is discussed in more detail in Section 1.3 below. We will refer to the resource with which this report is concerned as the “routing-addressing space” rather than IP address resources, while recognizing that addresses and routes are distinguishable parts of that resource space.

1.2. Common pool regimes

Various discussions of IP addresses, including statements by the RIRs, assert that addresses are “public resources.” But “public resource” is an unscientific term that can mean many different things to different people.[2] To add more precision, economic theory makes a useful distinction between four broad classes of goods: public goods, private goods, club goods, and common pool resources. The distinctions hinge on the degree to which resources are *rival in consumption* (i.e., one person’s consumption does not prevent anyone else from using it) and *exclusive* (i.e., the degree to which an owner or appropriator of the resource can prevent others from appropriating it.) These categories provide a more accurate and useful framework for the analysis of resource management regimes.

Table 1

[2] It might mean state-owned; or privately owned but state-regulated; or publicly shared; or a public good in the strict economic definition described below, or an essential facility, etc.

Classification of goods (adapted from Ostrom, 2005, p. 24)

		Subtractability of use (Rival occupation or consumption)	
Difficulty of excluding users		*Low*	*High*
	Low	**Toll goods**	**Private goods**
	High	**Public goods**	**Common pool goods**

With public goods, consumption is nonrival and exclusion is difficult or impossible. With private goods, consumption is rival and exclusion is relatively easy. Thus for private goods, market allocation (usually shaped by government regulations or subsidies in some way) is the most common option; the resource is allocated by means of exchanges of property rights among private owners. Prices will go up as supply goes down and conservation incentives will adapt accordingly.

Common pool resource management regimes are responses to a unique set of economic conditions. Consumption is rival, and thus private appropriation must be rationed or limited in some way. But for common pool resources there are factors that make it inefficient or undesirable to rely on exchanges of exclusive property rights to allocate the resource. It may be that exclusion is too costly, or that the interdependencies among users of the resource make it too difficult to define and enforce tradable property rights, or it may be that other forms of limiting appropriation are more efficient than market exchanges. An example of a resource which faces difficulty of exclusion is schools of fish in the ocean, which cannot easily be fenced in. In these cases collective governance rules can take the place of market prices as the allocator of the resource.

Resources do not fall unambiguously into these categories, and their status can change. Technological change, for example, made it possible to exclude owners of a radio receiver from access to a broadcast signal through encryption, transforming it from a public good situation to a private good.

1.2.1. Rival consumption

For IP address resources, usage is rival or (in Ostrom's terms), subtractable. Address assignments or allocations are rival because they must be exclusive and globally unique to function properly. If one network occupies an address block, another network cannot also use it on the global Internet, because simultaneous use of the same addresses would create interference or conflicts in routing. Exclusion, on the other hand, presents a less clear picture.

1.2.2. Exclusion

When a user appropriates a natural resource like an agricultural plot or fish from the ocean, the mere fact of possession by one person physically prevents others from also appropriating the resource. Exclusion of others from use of a virtual resource like numbers is not as straightforward, however. The act of assigning numbers to host computers on a network does not by itself prevent anyone else from also assigning those numbers to their own hosts. To maintain the exclusivity of assignments requires collective action in the form of a registry accepted by IP network operators as an authoritative coordination instrument.

IP address registries meet this need, keeping track of which organizations are using which address blocks and ensuring that the allocations are exclusive; i.e., that no two networks are given the same blocks. The registry's ability to maintain exclusivity, however, depends heavily on its acceptance and recognition by Internet service providers as a guide to their routing decisions. To enforce exclusive assignments of IP addresses, ISPs must refuse to route traffic to another network's address if it is not registered as the legitimate holder of that address. Here again we see the interdependence of addressing and routing.

Once one understands how registries contribute to exclusivity, it becomes clear that it is perfectly feasible to have exclusivity in the occupation of IP addresses. Laws, regulations or contracts could give registered address block holders a right to take legal action against "trespassers" who occupied address blocks that conflicted with those of a registered user. Thus, the choice of a common pool regime and its restrictions on address ownership and trading were *not* motivated by difficulties of exclusion. Rather, common pool governance emerged because of problems associated with limiting appropriation (see section 1.2.3 below) and issues associated with scalable routing (see section 1.3.2 below).

1.2.3. Need for appropriation limits

The absence of trading and market prices in a common pool regime means that some other mechanism must be used to limit appropriation of the resource. In the absence of tradable property rights, common pool regimes conserve the resource by collectively establishing and enforcing limits on appropriation. In IP address management, it is the address registry that acts as a gatekeeper to the address space, imposing administrative limits on the allocations or assignments any appropriator can claim for itself. In practice, these limits are based on a variety of policy factors, such as engineering-based assessments of a network's "need" for addresses, utilization levels and so on.

As noted before, given an authoritative registry that is respected by a critical mass of the world's ISPs, and a legal framework that prevents one party from occupying or using address blocks that have already been registered by others, exclusion is perfectly feasible. If we look more closely at the theory and practice of IP address management under the RIRs, we discover that the choice of common pool governance involves the need for appropriation limits. Current managers of the address space oppose rationing of the resource by price, fearing that it would lead to private property rights in addresses (which might have a deleterious impact on the efficiency of routing) and open the door to speculation, hoarding and other forms of unproductive occupation of address blocks. Because of these concerns, they prefer to set appropriation limits based on administratively-established technical criteria. The object of such procedures is to ensure that applicants for address resources will actually put them to productive use in real networks and not simply occupy them for their resale or pre-emptive value. These procedures are known in the industry as "justified needs assessments." Applicants for address space literally submit network plans and information about utilization levels; the RIRs review these plans and award address blocks accordingly.

1.2.4. Reclamation and reuse

In a common pool regime governing a fixed resource, the issue of resource reclamation is an important and rather neglected topic. Those who have been assigned or allocated addresses but no longer use them are, in principle, supposed to return them to the common pool to make them available for use by others. This potential for reuse is one of the primary efficiency benefits of a common pool regime. Resource reclamation has been a persistent weak point of the current address management regime, however. There are few positive incentives to return unused blocks. RIRs have no contractual authority to reclaim IPv4 resources from so-called “legacy” holders who received their allocations before the contractual regime was put into place. Even for organizations that contract with the RIRs, the ability to monitor the actual use of address and reclaim resources from current holders is weak. The scalability of detailed monitoring is questionable.

1.3. Routing and addressing

Routing is the automated process that directs Internet protocol packets from their origin to their destination. On the Internet, address allocation structures and routing protocols and standards are technically interdependent. The scaling properties of routing and routing equipment are affected by the way address blocks are distributed among networks. In order to understand these interdependencies, we have to back up and describe some basic Internet routing principles.

1.3.1. Routing tables

IP addresses can be described as part of the language that routers speak to each other. Internet routing protocols consider the IP address to be composed of two parts: the address of the network (the *prefix*) and the address of the connected computer (the *host*). Routing through the Internet is based on the network portion of the address. A router stores its best and alternate routes for each prefix in a routing table and uses this information to construct a forwarding table that controls the movement of each incoming packet to the next hop in its journey. Routers also transmit announcements to other routers about the address prefixes to which it is able to connect, and this information is incorporated into the tables of other routers. Thus, routers are engaged in constant, automated conversations with each other that exchange network prefixes and other information to keep every router informed about how to reach hundreds of thousands of other networks on the Internet. Currently, interactions among routers are based on a protocol known as BGP.

1.3.2. Externalities in inter-domain routing

These exchanges among independent network operators are rife with externalities. By externalities we mean that the behavior of one market actor can impose costs or benefits on the other without any voluntary transaction between the two that would enable the distribution of costs or benefits to be negotiated. As Huston notes, “the BGP routing space is simultaneously everyone’s problem, because it impacts the stability and viability of the entire Internet, and no one’s problem, in that no single entity can be considered [responsible for] manag[ing] this common resource.” (Huston, 2001, p. 13)

There are two major routing externalities: 1) the rate of change in one network’s routing announcements, and 2) the size of the routing table. Rapid change in routing announcements can increase the processing load of other network routers all over the world. And if no policy limits are placed on the number of route advertisements, it is possible that the size of the routing table in the core of the Internet (or default-free zone) could grow until it exceeds the processing power of the Internet service providers’ routing equipment. This is an externality because when one network adds announcements of many different fragments of an address block to the routing space it does not make its own operations much more expensive, but when such behavior is repeated across many other actors it makes the size of the routing table used by all ISPs larger and larger, making routing equipment more expensive. Some have described this problem as a tragedy of the commons.[3]

During the early-mid 1990s, as the Internet took off as a public medium, the number of prefixes listed in routing tables began to grow at an alarming pace; some felt that the scaling problem threatened the viability of the Internet. In an attempt to control this problem, the IETF and ISPs introduced provider-based address aggregation. Provider-based address aggregation is a hierarchical approach to address allocation that minimizes or eliminates direct assignment of IP addresses to end users. It gives network service providers (ISPs) larger address blocks from the registry, and ISPs announce this entire block into the exterior routing domain. Customers of the provider use sub-allocations from this address block, which are aggregated by the provider into a single route announcement and not directly passed into the exterior routing domain. This minimizes the number of entries in the routing tables, and also reduces the amount of traffic exchanged among routers – two very important economic efficiency benefits.

But provider-based route aggregation has two other economic consequences. First, it prevents IP address blocks from being portable across providers, which increases end user switching costs in the market for internet services. Second, it militates against trading or other uncontrolled transfers and deaggregation of address blocks among end users and organizations, which reduces the efficiency of address space utilization. The "assignment of addresses to hosts must follow the connection topology of the network in order for hierarchical aggregation to be successful." (Rekhter, Resnick and Bellovin, 1996)

---

[3] Cowie, 2009, Huston, 2001.

2. The techno-economic features of IPv6

Following on the general discussion above, this section focuses on the IPv6 routing-addressing space specifically. It analyzes the standards documents (RFCs), fee structures and policies pertaining to IPv6, and shows how certain assumptions about scarcity and allocation policy are embodied in them.

2.1. IPv6 and conservation

With 2^{128} addresses in the IPv6 space compared to the 2^{32} addresses of IPv4, the new protocol constitutes an enormous expansion. Mathematically, the IPv6 address space is almost infinitely large. This has encouraged some uncritical complacency about allocation policy. A European Union report, for example, states that "IPv6…would provide more locations in cyberspace than grains of sand on the world's beaches."[4] It is also commonly claimed that there are more IPv6 addresses than there are "atoms in the universe." However, a closer examination reveals that concerns about conservation still exist and cannot wisely be ignored.

2.1.1. Allocation policy for ipv6

The groundwork for IPv6 address allocation policy was established by the IETF, through its executive-level committees IAB and IESG. The IETF has made two important recommendations.

First, it released only 15 percent of the available global unicast IPv6 pool for allocation (a /3 block), and left the remaining 85 percent "reserved." The reservation was

[4] Europe's Information Society thematic portal, "IPv6: Enabling the Information Society." http://ec.europa.eu/information_society/policy/ipv6/index_en.htm

explicitly acknowledged as a hedge against the possible need for different, possibly more conservative allocation policies in the future.

Second, the IETF upped the scale of allocation policy. The critical principle underlying IPv6 allocation policy is that allocations are made to applicants based not on the number of individual IP addresses in a block, but on the number of *subnets*.[5] Subnets are large blocks of addresses which can be used to form networks. The smallest unit of subnet allocation in IPv6 is the /64, a subnet size that contains the capacity for 18.4 thousand trillion (18,446,744,073,709,500,000) individual IPv6 addresses. It is contemplated that /64 blocks will be assigned to home users or mobile phones. This unit was chosen, in part, so that it could incorporate an Ethernet physical address (EUI-64).

The basic unit for making allocations to organizations or end user sites is the /48. A /48 subnet contains 65,536 subnets of the /64 size. The basic minimal unit for Internet service providers is supposed to be the /32, which contains 65,536 /48 subnets and 4.3 billion /64 subnets. In other words, a /32 contains as many /64 subnets as there are addresses in the entire IPv4 address space.

These decisions are set out most clearly in two key RFC documents: RFC 3177 (2001), “IAB/IESG Recommendations on IPv6 Address Allocations to Sites;” and RFC 5375 (2008), “IPv6 Unicast Address Assignment Considerations.” RFC 5375 openly notes the detachment of allocation policy from host counts: “The practically unlimited size of an IPv6 subnet (2^{64} bits) reduces the requirement to size subnets to device counts for the purposes of (IPv4) address conservation.” Likewise, measures of “host density” for IPv6 refer not to the actual number of computers or other devices connected, but to

---

[5] These decisions were rooted in RFC2374 and RFC2450, where the IETF's IPNG working group recommended that the address block given to a single edge network which may be recursively subnetted be a 48-bit prefix.

the number of different network “sites” to which a subnet has been assigned: “…with IPv6, [host density] is calculated for sites (e.g., on a basis of /56), instead of for addresses as with IPv4.” (RFC 5375)

While its approach may seem extraordinarily liberal, RFC 3177 enumerated several reasons, many of them attractive, for taking this approach to allocation:

- A standard, provider-independent boundary between the network portions of address allocations is supposed to make it easier for users to change ISPs without internal restructuring or consolidation of subnets.[6] Even without a change of ISPs, these standardized boundaries make renumbering easier.
- Large initial allocations minimize the cost burden and administrative overhead associated with “needs assessment” by the RIR; it allows ISPs and their subscribers to grow substantially without any need to return to their ISP or RIR with additional (possibly costly, and possibly non-contiguous) requests. This is, in fact, a very important policy departure from the RIRs’ IPv4-era policy of increasingly restrictive and bureaucratic needs assessments. Initial allocations would no longer be based on “demonstrated need” but on some kind of basic classification of the applicant.
- This approach also allows sites to maintain a single reverse-DNS zone covering all prefixes

In their implementation of these guidelines, most RIRs have proposed /32 as the basic initial unit of allocation to Internet service providers. A /32 would give ISPs 65,536 /48s to allocate or assign to their customers. Before requesting more, the RIRs initially

[6] For example, if ISPs consistently hand out /48s to end user organization sites, regardless of size, then an organization that switches ISPs can renumber easily by simply changing the network prefix in all their site addresses. (Whether ISPs will go along with this remains to be seen.)

proposed to require that these blocks attain a Host-Density ratio of .80. Later, the recommended HD ratio was increased to .94.

2.1.2. Does conservation even matter?

If the IPv6 address space is so indescribably large, do we need to worry about conservation at all? As the discussion above explained, the basic units of allocation are also extremely large, and the distribution of such blocks will, without question, result in the "waste" of vast quantities of bit combinations that could in theory be used as addresses. In the intermediate term, at least, very few home networks or small office assignees of /48 address blocks are likely to make use of any but a tiny fraction of the 1,208,925,819,614,620,000,000,000 bit combinations their subnet contains. Some may grow into that space or discover new applications of networking that make use of it, but many will not. Still, in RFC 3177 the IAB claims "we feel comfortable about the prospect of allocating 178 billions /48 prefixes under that scheme before problems start to appear. To understand how big that number is, one has to compare 178 billion to 10 billion, which is the projected population on earth in year 2050." (RFC 3177) In an attempt to raise a yellow flag of caution, Thomas Narten (2005) noted that giving every human on the planet a /48 would occupy "fully 1% of the available address space…in 50 years," which he claimed showed that the address space was "nowhere near practically infinite."

While the IETF statement is optimistic and expansive and the Narten statement is cautious, both calculations reveal how little attention is paid in Internet circles to the economic importance of *reclaiming* abandoned, unused or underutilized addresses. Both statements assume that giving a /48 address block to each member of the whole population on earth is a static, one-time decision that would result in an outflow of only

10 billion /48 address blocks after 50 years. But this calculation is overly simplistic and, as a result, far too low. We have to calculate the ongoing *flow* of address blocks to a population that is constantly adding new members (hence making new allocations) and whose ranks are also depleted by deaths (and therefore abandoning allocations which must either be reclaimed or wasted). Assume for the sake of argument that none of the allocations freed up by deaths are reclaimed. In that case, depending on the birth and death rates, the number of address block allocations required to give every member of the human race a /48 increases by a factor of three to four – it could consume 30 – 45 billion address blocks, not just 10 billion.[7] Note that organizational populations have births and deaths, too; that is, ISPs, other service providers, and end user networks are constantly coming into and going out of existence. The effective utilization rate projected must be increased accordingly.

Uncertainty about future use, and historical precedent, creates another reason for caution about the need for conservation. The IPv4 space was also considered to be enormous and practically inexhaustible in the early days of the Internet's development. For the first decade of the Internet's existence (1982 – 1992), class-based allocations were made to government agencies and private sector organizations participating in U.S. government-funded research and development.[8] Near the end of this period as the Internet gradually opened to larger sectors of society and the world, a surprisingly large amount of the IPv4 address space was allocated. Table 2.1 displays the numbers.

[7] One must sum the initial allocation to the population, all births, and all deaths over the period.
[8] The earliest address allocations were based on the three classes of addresses defined in RFC 791, the basic document defining the Internet protocol. Class A addresses, which correspond to what we would now call a /8, have a network part of the address that is 8 bits long, leaving room for unique addresses for 16.7 million hosts. Class B addresses (now referred to as a /16) had a network prefix of 16 bits, leaving room for 65,556 unique addresses. Class C blocks (a /24) allowed for only 256 unique addresses for hosts.

Table 2
Network Number Statistics as of May 1992
(Source: RFC 1466)

	Available	Allocated	Percent allocated
Class A	126	49	39%
Class B	16,383	7,354	45%
Class C	2,097,151	44,014	2%

It is apparent from these statistics that the initial allocation policy embedded in RFC 791 and its implementation not only underestimated overall demand, but completely failed to anticipate the actual *structure* of demand. Demand for address blocks was concentrated in the Class B range, whereas demand for small, 256-host networks was much lower than expected. Moreover, classful allocations were structurally wasteful, in that applicants whose need fell somewhere between the sizes of the defined classes had to be given the next highest block size regardless. From an economic standpoint, there are structural similarities between classful allocations and the IPv6 recommendations of the IETF.

The impact of these legacy allocations still lives with us today. Latecomers to the Internet party have been greeted with more restrictive address policies. A document submitted to APNIC in 2005 by Millet and Huston called attention to the possibility of a more rapid than expected occupation of available address resources due to the capacious initial subnet allocation policies. Millet and Huston criticized the IETF claim that "if the [IETF] analysis does one day turn out to be wrong, our successors will still have the option of imposing much more restrictive allocation policies on the remaining 85%." "From a public policy perspective," Millet and Huston wrote, "we stand the risk of, yet again, visibly creating an early adopter reward and a corresponding late adopter set of barriers and penalties." This campaign from leading Internet figures in 2005 led to a tightening of RIR initial allocation policies. Millet and Huston proposed (and ARIN and APNIC eventually adopted) a more conservative policy on initial IPv6 allocations, offering a /56 rather than a /48 to certain classes of users, and increasing the HD ratio required for additional allocations from .80 to .94. The latter change, according to Huston's calculations, would reduce address occupation by a factor of 10 over a period of three years.

What is clear from this analysis is that conservation or appropriation limits of some kind are still needed, but that we can indeed afford to be much more liberal with initial allocations. A critical factor in conservation will be efficient reclamation policies and incentives. Neither Narten, nor Millet and Huston, nor the IETF discuss the incentive for occupiers to release unused or underutilized address resources. Providing incentives for more efficient reclamation could extend the life of the usable address space considerably.

On the other hand one must not underestimate the costs associated with withholding addresses from enterprises unnecessarily. A century is a very long time for a communication protocol to last. It is difficult to think of a single electronic communication system that has survived that long without undergoing a fundamental transformation. Withholding addresses that could be profitably used on behalf of future occupants or applications that may never materialize is also a form of waste.

The main issue here is one of uncertainty about where is the optimal point on a tradeoff curve. This is exactly where economic incentives might be most useful as a mechanism for adaptation.

2.1.3. RIR fees and IPv6

Address fees are a potentially important conservation tool. Officially, the RIRs present themselves as member-based organizations, and their fees as membership dues that recover the cost of their services. Most of them are uncomfortable with any assertion that they are charging for IP addresses. The RIRs do provide important services, such as maintaining the registry and the Whois lookup, performing “needs” assessments, supporting policy-making processes, and so on. On the other hand, some RIR fees reflect a positive correlation between the fee size and the size of a member’s IP address allocations. And it also seems logical, in both equity and efficiency terms, that people who occupy more address space should pay more, especially when they are commercially exploiting the addresses. But there is an important difference between fees as methods of recovering the cost of RIR services and fees as rationing devices to encourage conservation of the address space. The RIRs are currently sitting in an uncomfortable space between the two.

The researchers compiled detailed information about the fee structures of each RIR for IPv6 allocations, which is available upon request. This data yields five observations.

First, insofar as currency differences make direct comparison possible, there is minimal variance in the fee levels across the RIRs for address blocks of size /32 and smaller. The normal fees for /32s among ARIN, LACNIC and AFRINIC are identical (and are all denominated in USD). The fees for address blocks larger than /32, however, vary widely across RIRs. RIPE-NCC seems to have slightly lower rates (denominated in Euros) and APNIC seems to have higher rates (denominated in Australian dollars).

Second, the fees charged by three RIRs for IPv6 addresses (ARIN, APNIC and RIPE) tend to get larger as more addresses are occupied, but only above the /32 level. This implies that fees are intended to perform a conservation function, and these conservation incentives seem to be directed at larger ISPs. Fees increase most slowly with size at RIPE-NCC, and more rapidly at ARIN. AFNIC and LACNIC, on the other hand, have very simple, two-part fee structures. They only distinguish between “small” and “large” allocations, with the dividing line being the /32. All RIRs seem to discriminate between recipients of /32s and /48s based on technical assessments of need and do not seem to rely much on pricing to encourage conservation of address blocks at that size. However, a recent fee restructuring at APNIC ties fees to block size directly. APNIC fees respond logarithmically to increases in block size. This policy is too recent to assess, but could constitute an interesting new trend in RIR resource management.

Third, and contributing further to the view that fees serve a policy function, ARIN has discounted their standard IPv6 address fees from 2008 to 2011 in an attempt to encourage migration to IPv6. Post-discount, the initial fee for a /32 allocation from ARIN in 2009 is a paltry $562.50. The discounted ARIN fees are very similar in size and structure to the RIPE-NCC fees. Similarly, LACNIC has created an exception for IPv6 fee schedule; Members having only IPv6 addresses will continue to be exempt from paying membership fees until July 1st, 2012. -- If an ISP requests only IPv6, there is no charge in 2009 to 2011; however, after 2011 the ISP need to pay the regular fee based on the current fee schedule. As far as we can tell, no RIR has raised fees for IPv4 addresses as the free pool approaches depletion.

Chart 1: Comparative fee structures for IPv6 addresses, the 5 RIRs

(10-year period, 1.00 EUR = 1.39 USD; 1.00 AUD = 0.792 USD)

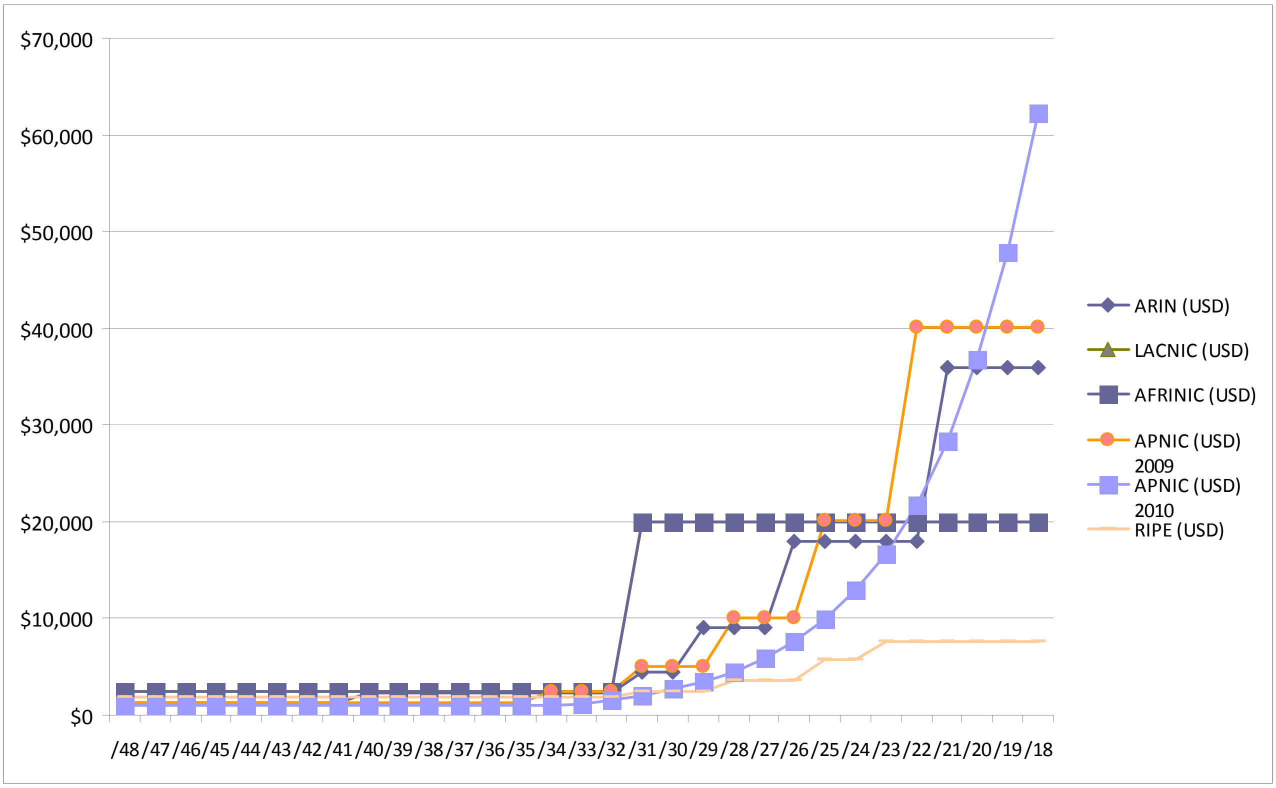

Chart 2: ARIN IPv6 fees (10 year period)

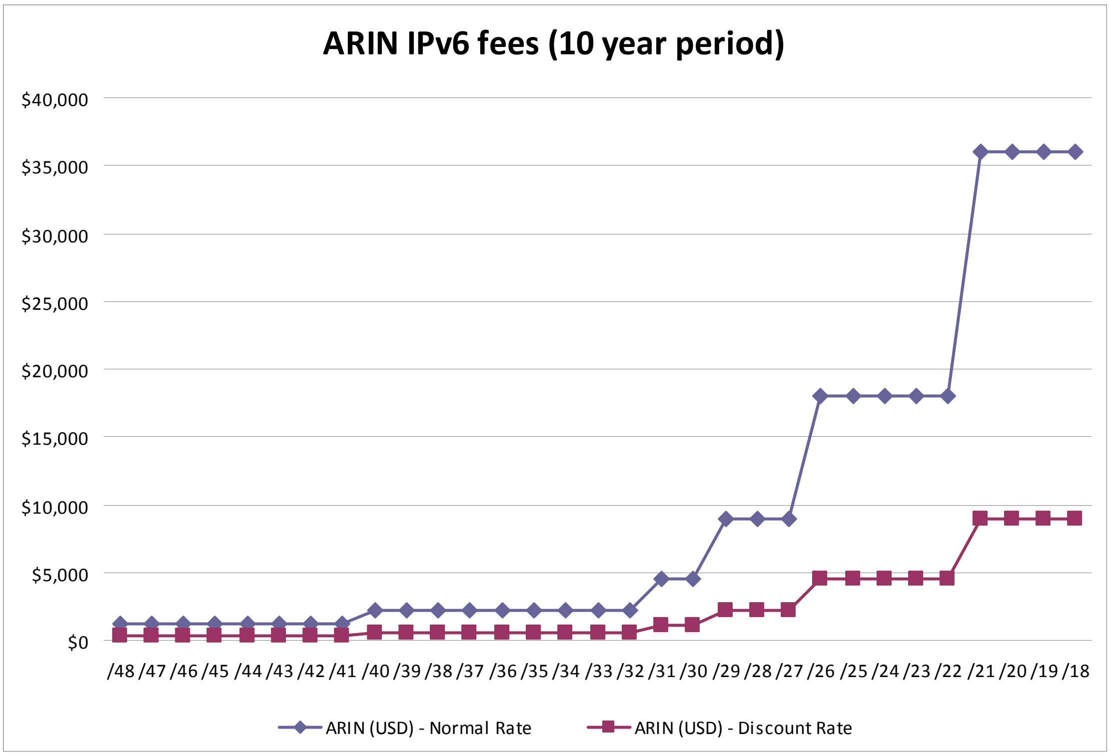

Fourth, if one considers the fee to be related to the amount of address space consumed, RIR fee structures provide massive volume discounts to ISPs (with the exception of APNIC's recently-reformed fee structure). The current fee structures of AFRINIC and LACNIC, for example, value a /32 address block at about $2,500 per year. An Internet service provider that gets a /32 block from an RIR at that price is paying about $0.04 per year for each of the 65,500 /48 address blocks it contains. On the other hand, an organization site that is awarded a /48 is likely to pay $2,500 a year for it as well. If the larger ISP asking for a /32 were charged at the same rate as the small organization for all the /48s subnets in its /32 block, the ISP would pay $163.8 million a year for a /32. Nominally, RIR fees are based on the premise that ISPs are entitled to a fixed amount of IP addresses if they have a certain number of customers (IPv4) or end sites (IPv6). However, the RIRs also try to charge larger providers higher annual membership fees (which is not unusual for trade associations), and it calibrates the "size" of its member organizations in terms of the number of IP addresses it occupies. The overall effect is that larger occupiers of the space do pay more, but they also enjoy an enormous volume discount in the IPv6 space (about 10^5). Also, while there is much noise against the idea of "paying for address space" in some quarters, the fact of the matter is that ISPs use address space as an input into their services, and then turn around and sell service to their customers, often charging per IP address.

Fifth, the experience of APNIC with National Internet Registries (NIR) shows that RIR "membership fees" do rely on access to addresses as one of the key "value" components of the payment. Under an NIR structure, it was possible for many ISPs to join together and pay a single membership fee to APNIC. This would allow many ISPs to jointly pay a single membership fee, and gain access to a potentially large number of addresses. Fearing exploitation of this, APNIC for a time imposed "per address fees" on NIRs. One can understand APNIC's dilemma. As an extreme example, suppose that all of a region's ISPs joined together into one club and paid one $36,000 membership fee to its RIR. If addresses were then allocated solely on the basis of "need," no RIR could sustain itself. It should be obvious that fee must be tied to addresses in some manner.

2.2. IPv6 and routing economics

Initial policy toward IPv6 allocations was heavily influenced by the scaling problems associated with routing that were experienced in the IPv4 space. This is one reason why the initial allocations were proposed to be so large. Large initial allocations would allow

ISPs and end-user organizations to announce one network prefix for a long period of time, rather than expanding through acquisition of additional blocks that might require additional announcements.

On the other hand, the larger address size of the IPv6 space means that the routing tables will have to carry bigger address prefixes. This will increase the demand on the memory of routers. And inherently associated with the vast number of subnets is a potentially much larger universe of routes. This has given some experts concern. An IETF report noted, "Given that the IPv6 routing architecture is the same as the IPv4 architecture (with substantially larger address space), if/when IPv6 becomes widely deployed, it is natural to predict that routing table growth for IPv6 will only exacerbate the situation." Another sobering fact is that during the transition period from IPv4 to IPv6, routers will have to carry routing tables for both protocols.

Until adoption increases, it is difficult to know whether, or for how long, the larger initial allocations of IPv6 will outweigh the other factors. But in pure potential, the vast size of the address space and the vast number of subnets and hosts that could be connected through it means that the scaling problems of 1990-1993 could pale in comparison.

There are conflicting opinions about the routing problem within the technical community. Huston (2009) claims that growth in the routing table is predictable and falls within the limits of Moore's Law, which will bring increases in the processing power of routers that compensate for growth in routing table size. Huston concludes that the life of BGP can be extended for the intermediate future, at least. On the other extreme, there are those who believe that BGP is inherently fragile and that an adequate response to it scaling problems, along with its security issues, requires replacing it with a new routing architecture that separates locators from identifiers. Ironically, this locator-identifier split would mirror the kind of routing that the mobile phone system currently employs. There are several proposals to do this,[9] none of which has been fully developed or has broad consensus.

It is outside the scope of this report to project the future of Internet routing or to make specific proposals regarding which routing structure the Internet should adopt. Given the interdependence of addressing and routing economics, however, any proposals for a market system in address allocation needs to keep the uncertainty surrounding this issue in mind.

3. Introducing economic factors

This section makes a specific proposal for how the RIRs could move to a more market-oriented method of address allocation and route aggregation. The proposal is intended to create a functioning market for address blocks in a gradual manner. It would offer a new option to ISPs and end user organizations seeking addresses from RIRs, a "transferable address block lease" (TABL).

3.1. The initial allocations problem

The first problem facing any market-oriented approach to the IPv6 space is how claimants initially occupy IPv6 address blocks. Before one can have a market price system one must have exchanges or transfers among market participants who possess title

[9] This includes O 'Dells GSE proposal, 8+8 Addressing Architecture, ILNP (Identifier, Locator Network Protocol), HIP, LISP (Locator/Identifier Separation Protocol); FARA (Forwarding Directive, Association and Rendezvous Architecture); and ISLAY - a new routing & addressing architecture.

to parts of the resource. Since the vast majority of the space has not yet been occupied, what rules or processes would govern the first appropriations of the IPv6 space by private users? These initial allocations obviously cannot be purchased for a market price, because there is no market yet. But one cannot simply allow anyone to claim as much as they like, because that could lead to a land rush and wasteful or unfair forms of over-occupation.

Let us refer to that as the *initial allocation problem*. Under present arrangements, initial allocation is handled by means of administrative and technical criteria established by RIRs, as described above. Any attempt to move to a market system needs to either build on that, or create an alternative method of establishing and recognizing initial claims.

3.2. Auctions?

Auctions were considered, but are *not recommended*, as a mechanism for making initial allocations. The virtue of auctions is that they establish a "real" price for the resource. Applicants would express their valuation of some portion of the address space through bids. Competitive bidding among applicants would approximate a market price for some portion of the resource, and award it to the organization or entity that valued it most highly. But the value of auctions in this case would be undermined in this instance by the need for top-heavy implementation and by deep uncertainties about how and when to hold an auction.

It is immediately apparent that one cannot have auctions without an auctioneer. The auctioneer has to define the resource units to be auctioned, establish the process or method of bidding, administer the auction fairly, and then make authoritative transfers of portions of the resource to the winning bidders. The only feasible auctioneer in the near term would be the IANA or the RIRs. However, these organizations are based almost exclusively on technical rather than economic expertise. It is difficult to imagine a wholesale switch from their current methods to auction methods taking place in the near term.

There are other, more serious problems. Consider first the timing and scope of the auction. The address manager could either auction off the entire IPv6 address space at one time, or auction off portions of it at periodic intervals. Doing it all at once would not be realistic. The selection of any specific time for such an all-embracing auction would be arbitrary, and yet would have a major impact on IPv6 resource use for decades to come. A one-off auction would exclude that (very large) part of the world which is not ready or able to make use of v6 resources, as well as future generations. Moreover, since the value or degree of scarcity associated with the address space is still highly uncertain, the outcome would have very little use as a guide to the pricing of the resource. It is possible that all current bids could be satisfied at zero price, and an auction of the whole would become indistinguishable from an unrestricted land rush.

Holding auctions for smaller portions of the address space is more sensible, but creates its own set of problems. Breaking down the auctioning entities into regions might address some of the problems associated with unequal readiness in different parts of the world, and because the RIRs are already regionally structured this would not be difficult. The resource steward, however, would still have to make economically arbitrary decisions about the size of the address space to be auctioned off in any given period and how often to hold the auctions. Confining the auction's scope to too small a portion of the address space would drive up prices unnecessarily, as has happened with spectrum auctions. The manager of the auction, if it benefited from the proceeds, might also have an incentive to structure the auction to maximize revenue rather than social efficiency (as has happened with both spectrum auctions and telecommunication privatizations). Thus, at worst the price obtained via auctions would almost be as arbitrary as setting an administered fee. However, a sequential auction procedure of portions of the space is an option at least worth considering.

Another issue centers on *what* would be auctioned. The administrator would have to define the units of the resource to be auctioned off, and decide on the distribution of the different unit sizes. These decisions would greatly affect the bidding. Larger blocks would attract larger bidders; smaller ones would make it easier for smaller players to win but might affect the efficiency of routing. An address space manager making decisions about this would be thrust back into the position of trying to predict exactly which types of blocks were needed in which quantities. These *a priori* decisions, made in ignorance of actual market conditions years ahead, might have too much influence over the structure of the address space, just as the initial class structure of IPv4 addresses did.[10] If it is not costly to disaggregate and re-aggregate blocks as needed through future trades, this may not be a fatal problem. But if it created legacy structures, it might become a problem. For example, if address blocks are auctioned off only in /32 and /48 sizes, it might be more difficult and costly for those needing a /20 to acquire the contiguous blocks they need in the future.

All of the problems cited in this paragraph could be addressed through appropriate auction design. But that would mean a lot of effort would have to be put into the design

[10] If we take spectrum auctions as our guide, we see that the communications regulator plays a very large role in defining channel size and geographic scope of the spectrum units, and other technical factors such as equipment standards, interference standards and so on.

and administration of the auction process. And that is probably the most important argument against auctions as a system of initial allocation. Auctions could increase, not reduce, the bureaucracy and policy overhead associated with initial allocations.

Auctions work best when a fixed asset or well-bounded resource space has a number of well-known alternative uses and/or competing users, and there is already an established distribution of the resource. That is currently true of the IPv4 space. It is almost completely occupied and its alternative uses and users are well-established. Thus in IPv4 a competitive bidding process for the remaining allocations would serve a useful function. But that is not true of the IPv6 address space. It is not clear whether any current uses are truly mutually exclusive; it is possible that all current users can be accommodated by the large size of the space in the short term; too little is known about the demand for various block sizes; too little could be known about future demand. Therefore, if there is to be a market for IPv6 addresses, it should emerge gradually, from trades among parties who have already occupied IPv6 blocks, or some other method.

3.3. The Transferable Address Block Lease (TABL)

The proposal described in this section would employ two economic techniques for address block allocation: 1) fees related to the size of the blocks and 2) transferability. It would not create any disruption in policy or practice, but simply offer ISPs and end user organizations a new option, which they could choose to use when seeking addresses from RIRs if it seemed attractive.

RIRs would continue to field applications for address blocks under their established methods. But they would also set aside a block of contiguous IPv6 addresses, perhaps a /16, for a new form of allocation known as *transferable address block leases*

(TABLs). For these blocks of addresses, RIRs would set annual fees that increase with the size of the address block (see the section below for more detail on the nature and structure of the fees), and allow applicants to lease address blocks anywhere between the range of /32 and /48, inclusive.

While this would initially appear to be a more expensive option than obtaining addresses in the established way, it would have three compensating advantages. First, no needs assessment would be conducted for applicants wishing to lease these blocks. Second, in the TABL block, the registries do not distinguish between Provider-Independent and Provider-Aggregated address space. Third, the lease would be tradable; that is, it could be sold to another organization willing to assume the same contractual conditions attached to it. The block could only be transferred or sold as a whole; it could not be deaggregated and transferred in portions smaller than that.

It is proposed that the fees associated with these transferable address block leases would be slightly higher than the fees associated with the “application” and “membership” fees organizations normally pay for address blocks. The higher price for the transferable lease option would reflect the stronger rights over the address block, and hence greater value, leased to the applicant. The holder of the lease would retain exclusivity over the block as long as they paid the fees. Contractually, it would be made clear that the blocks are leased, not owned in fee simple. The right to exclusive use of the addresses, or the right to transfer them to other parties, would be contingent upon the holder’s compliance with contractual conditions deemed essential to maintaining the integrity of the registry function and other rules of good conduct.

As a precautionary measure, we propose to set limits on the number of transferable blocks any single organization could lease (e.g., no more than four /32s). If demand for the blocks allocated in this fashion depleted their availability more rapidly than expected, the RIRs could respond in one of two ways. First, they could simply expand the TABL block, adding another /16. Or they could begin to hold auctions for them rather than relying on fixed fees, using the prior indicators of demand to guide the quantities and sizes made available for auction. If, on the other hand, the RIRs found few takers for these rights, they could either reduce the fees or some of the limitations associated with the transferable lease rights.[11] Or, they could call off the experiment altogether and simply stop issuing new TABLs.

One of the virtues of this proposal is that it is incremental and opt-in: the RIRs continue with their normal procedures. If they discovered that many applicants preferred the new TABL approach, it would provide them with valuable information about the value of address blocks and the problems associated with their legacy procedures. Longer term, such an approach would provide an entry into a full-fledged market system, allowing the RIRs to concentrate on the vital registry function and leave the determination of need to normal market processes.

3.3.1. Route aggregation and TABLs

The primary obstacle to the tradability of address blocks has always been the route aggregation problem. We address this problem in a simple and effective manner by fixing the size of TABL blocks at /48 at the bottom, and /32 at the top, and by limiting all trades to the original TABL units leased. In other words, an organization that leased a /48 could

[11] E.g., they could allow deaggreation of the blocks in units of /48 and higher.

only transfer the entire thing; an organization that held a /46 TABL could only transfer the entire /46. Organizations who discovered that they had too much address space and wished to reduce their fees could try to trade with someone holding a smaller TABL and seeking a larger one, or if unable to do that return it to the RIR and apply for a needs-based assessment.

The TABL option might be most attractive to medium-sized organizations and service providers seeking Provider Independent address blocks. Encouraging PI might also have implications for routing table growth. However, existing policy toward IPv6 allocations offers end user organizations /48s and even /56s. Thus, it is difficult to argue that the proposal would make the current system any worse or any better with respect to routing table growth.

4. Pros and Cons of the proposal

In this section we lay out the projected benefits and criticisms or concerns that might be leveled against the TABL proposal.

4.1. Benefits

The TABL approach has a number of benefits.

- It frees both the RIRs and the applicants from the complexities and constraints of needs assessment, providing a simpler and more controllable method for organizations to acquire address space.
- It gives end users more independence from providers and therefore increases competition in the ISP market

- It creates better incentives for conservation by calibrating fee sizes to block sizes more closely
- Through the use of recurring fees based on the size of the block, it incentivizes reclamation of address space that has already been allocated
- It counters some of the pressure that might lead to overly liberal allocations in the early years
- It can be implemented in a gradualistic, way, and if it fails the experiment can be stopped

The main advantages are explained in more detail below.

4.1.1. TABLs vs. Needs Assessment

The chief advantage of TABLs is that they make the most desirably-sized address blocks available to anyone who wants them in the quantity they (as opposed to a central planner) think they need, through a simple procedure. It sets a fixed fee for a certain amount of addresses and allows anyone willing to pay that fee to get the addresses. Organizations can opt out of needs assessment, although needs assessment remains in place for most of the address space. Table 3 summarizes the merits of fees vs. needs assessments.

Table 3:
Comparison of methods for limiting appropriation of IPv6 resources

	Method of limiting appropriation	
	Needs assessment	**Fees**
Strengths	Ensures productive use of addresses in initial allocation Cheaper for large ISPs and large users	Low administrative overhead Creates self-enforcing conservation incentive Creates self-enforcing reclamation incentive
Weaknesses	High administrative overhead No incentive to define "need" in ways that conserve No incentive to release resources after	If fee level too high, imposes needless cost If fee level too low, weak conservation and reclamation effects

	initial allocation Reclamation intrusive and expensive Questionable scalability HD calculations can be gamed	

Needs assessments force RIRs to apply very specific, often rigid technical and operational criteria to applications for address blocks. For IPv6, these criteria must be created in advance, in ignorance of all possible requests and use cases. It is likely that such an approach will have unintended consequences that prove to be costly obstacles to implementation in practice. Partly in response to that problem, the RIRs initial allocation policies for IPv6 are almost completely unrelated to “need” in the traditional sense. What they really are is a rough rule of thumb for defining the relative apportionment of addresses among different classes of users. These rules take the following simplified form: If you are an X, you qualify for a /48; if you are a Y, you qualify for a /32. The problem with that policy is that in the real world any given X may have special circumstances or valid reasons for claiming a /47 or a /46 instead of a /48, or the distinction between X and Y may erode over time or be blurry in certain instances. For example, if RIR allocations are based on the number of “sites,” then all kinds of games can be played with the definition of a “site,” which might either open the door to too liberal appropriation, or, conversely, leads to policies that are too restrictive. Any attempt to modify RIR policies takes a long time and significant investments in political activity, and offers no guarantee of success. In general, in an IPv6 world where both the number of addresses used and the diversity of applications increases, there are questions about the scalability of traditional needs assessment.

A TABL approach limits appropriation while removing the barriers associated with attempts to define administrative categories and rules that must be applied to a broad and diverse class of applicants. They let the applicant determine what size of an allocation they need; the resource administrator simply tries to charge them something approximating the social cost of excluding others. This removes some of the most important obstacles and uncertainties related to access to IPv6 addresses. The fees associated with different block sizes would be fixed and known, and those who use less would pay less. This could open the door to smaller players, innovators, and people with new ideas about addressing that might not make it through the traditional hurdles erected by the RIRs.

Would the lack of need assessments lead to widespread wasteful and inefficient use of the IPv6 space? We do not believe concerns about this are justified, for four reasons. First, there will be fees charged for TABLs, and the fees will be slightly higher than the normal ones, so people will not grab them for no reason. Second, TABLs will represent only a small portion of the available address space, and there are also limitations on the number that can be taken. The IPv6 address space is so vast that even if all of them are snapped up, and even if they are so popular that the IRs are pushed to expand the number available, it is a small portion of the overall resource space. Finally, as noted in Section 2.1.1, under existing policy most initial IPv6 allocations are not really based on "need" anyway, but on pre-defined categories.

4.1.2. End user control

TABLs are designed to eliminate the longstanding distinction between Provider Independent (PI) and Provider Aggregated (PA) address blocks. This could create the benefit of increased end user control over their network destiny, and reduce switching costs if they want to change providers. The current RIR regime, which favors aggregation of addresses under providers, assumes that ISPs will selflessly pass on the advantages of abundant address space to their customers. It also assumes that ISPs will go along with IETF recommendations about standardizing network-host boundaries and make it as easy as possible for customers to move from one ISP to another. This expectation is not realistic. Decades of experience with network operators in markets with high switching costs indicate that network operators will not make it easy for customers to switch. The TABL approach gives end user organizations stronger bargaining powers. If ISPs attempt to overprice addresses or use them as lock-in techniques, users have a viable alternative.

4.1.3. Conservation and reclamation

By tying the size of address fees more explicitly to the size of address blocks allocated, TABLs create a stronger incentive for conservation. And its use of recurring fees and the transferability option facilitate better reclamation of unused or underused resources. Occupiers of address blocks would be given strong monetary incentives to release blocks that they don’t need. Transferability also ensures that if scarcity of some kind does emerge in the long term, number resources will move to their higher-valued applications in a smooth manner.

4.1.4. Latecomers/fairness

As noted earlier, there are legitimate fears that IPv6 allocation will repeat the history of IPv4: a pattern of over-occupation in the early years, when it appears that the address space is unlimited, leading to massive waste and misallocation, followed by tightening of policy that imposes more restrictions and higher costs on latecomers. The combination of fees and transferability mitigates these concerns somewhat, but does not completely eliminate them. Fees + TABLs improve the situation in several ways. First, a needs-based allocation system inherently rewards existing, developed Internet economies. Large occupiers of address space would be able to get enormous IPv6 blocks virtually for free simply by showing that they can use them; undeveloped economies would not qualify for large allocations because they could not put them to immediate use. Second, fees provide an incentive for early adopters to mitigate their claims in line with conservation needs. In particular, the matching of fee size to resource quantity ensures that all early adopters have an incentive to release unneeded resources. Third, TABLs allow those willing and able to pay the leasing fee to reserve up to four /32 address blocks at any time, even if they are not yet able to use them. Fourth, transferability makes it possible to reallocate address blocks more quickly and flexibly as demand shifts.

However, if scarcity develops latecomers would probably pay higher costs to acquire resources than the early adopters. Thus, there could still be a latecomer disadvantage. If the policy objective is to ensure that all countries, *qua* countries, have absolutely equal rights to claim IPv6 resources then some kind of equity-based allocation principle would have to be used. Such an allocation, however, would create a massive sacrifice of economic efficiency, by reserving large amounts of resources to entities that could not use them, and (unless there is a TABL regime) might lead to the development of various kinds of gray or black markets for the underutilized resources.

4.2. Costs, questions associated with TABLs

Like any new idea, implementation of TABLs raises a number of concerns. Some of them are real concerns that need to be taken seriously; others are arguments that are likely to be made against it but lack merit and need to be dismissed. Here are some concerns:

- On what basis can the initial fees be set?

- Would the fees discourage IPv6 use?
- Where would the money go?
- Are there problems associated with a mixed system?
- Could someone “corner the market” for IPv6 address blocks?

4.2.1. How to set the fees?

The most important issue concerns how high or low the address block registry would set the initial allocation fees and establish the slope of the line at which fees would rise based on size. We have already noted the way in which the vast size of the IPv6 address space makes a purely linear fee structure prohibitive: organizations that now pay $2,500 for a /48 would, if this rate were extended linearly, pay $163.8 million a year for a /32. Obviously, that level of payment would defeat the purpose. A well-designed TABL fee might be lower at the lower end of the available address blocks (/48) and higher at the higher ends (/32). The registry must be sure to choose a fee structure that makes a TABL slightly more expensive than normal “free” needs-based allocations, but still attractive and feasible for those who could make profitable use of them. A major underestimation of fee sizes and too flat a rate structure might lead to a land rush; oversized fees could make the experiment fail, as no one would choose to make use of it.

While these are real issues, there is nothing inherently insoluble about this problem. The registries would need to carefully design fee structures, utilize simulations to test their effects, and make adjustments in their implementations. Because only a limited part of the available v6 address space would be set aside for TABLs, the risks of the experiment are contained.

4.2.2. Fees and IPv6 incentives

Some might claim that if TABL fees are big enough to have the desired effects of encouraging conservation and reclamation, they might act as an impediment to IPv6 adoption by smaller players; conversely, if the TABL fees are not big enough, they may not create sufficient incentive to conserve. It is true that pricing creates an economic cost and that this may be viewed by some as a barrier to IPv6 migration. There are two simple answers to this objection, however. First, under the proposal made here, no applicant has to pay TABL fees. They could continue to get addresses in the traditional way. Second, IPv6 addresses are not "free" even under the administrative system; fees are still charged to ISPs and end user organizations for "membership," and the preparation of applications for addresses is not costless. In that respect, this concern is not much different from the one expressed in the first bullet point: it is important that the fees not be either too large or too small. Administrative fees as a method of initial allocation are not really a radical departure from what the RIRs currently do. The main difference is that there should be a stronger relationship between the size of the address block occupied and the fee, and much less of an attempt to tie allocations to specific categories of users and uses.

4.2.3. Where would the money go?

This proposal consciously employs fees as tools of policy rather than as ways of recovering the cost of RIR services. In other words it uses fees to minimize social cost rather than to reflect RIR service costs. Although some of the revenues generated from such fees could be used to cover the costs of RIR services, they would bear no relationship to the cost of providing RIR services. If these fees generate an "embarrassment of riches" for the IR administering the TABL scheme what would we do with the funds? One response to this problem is straightforward: fees could be reduced if they generated too much money and there was not enough demand to worry about over-occupation of the resource. It is also possible to develop plans for trust funds that distribute the surplus to Internet development. We do not consider this a desirable option, however, because as a method of capital investment such programs tend to create undesirable dependencies and to reward grant-getting skills rather than self-sustaining forms of enterprise. This is an issue to be seriously concerned about, but again not an insoluble problem.

4.2.4. Can the market be cornered?

The charge that someone might be able to “buy up all the address space” via the proposed method can be dismissed readily. First, it would not be difficult to set and enforce a simple limit on the aggregate amount of TABLs a single entity could hold. We proposed such a limit in the TABL proposal (four /32s, or a /30). Second, even with low fee levels, tens of billions of dollars, if not trillions, in recurring annual costs would be required to maintain control of the entire IPv6 space. Neither large businesses nor major governments are likely to gamble with that kind of money. And with such large sums required, it would make no sense to even attempt such a strategy unless there was a reasonably certain prospect of controlling the entire IPv6 space. Third, as noted before the value of address blocks depends on largely voluntary decisions by Internet service providers to use a common address registry for routing purposes. Any rogue operator who succeeded in spending trillions to achieve such a position in the hopes of later charging monopoly rents to recoup its investment could be undermined by a decision by a critical mass of the world’s Internet operators to establish an alternative IP address registry and route around the monopoly.

4.3. Role of the ITU

The TABL proposal was initially conceived as a way of introducing economic incentives into the existing RIR regime. Indeed, it was designed in a way that deliberately minimizes altering the fundamentals of that regime; it is simply an option added to current methods. This does not, however, mean that the ITU could not play a constructive role in its implementation. There are two ways in which the ITU could get involved.

4.3.1. Model contractual/legal framework and fees

With its long-term development of the International Telecommunication Regulations, the ITU holds experience and expertise in international legal and regulatory frameworks. This capability could be put to use researching and developing a model framework for a TABL contract. The framework could recommend methods for dealing with such issues as: how to define the TABL right in contractual terms; conditions for revocation of the contract; conditions and procedures for transferring the TABL from one party to another while maintaining the Registry in the informational loop so that accurate records could be maintained; long-term guarantees regarding the fee size and protections against opportunistic increases; etc. In connection with this work, the ITU could also conduct research and simulation into the appropriate fees to charge for TABL blocks.

4.3.2. ITU as IR

The ITU could implement TABLs in its own address space. This option assumes, of course, that the ITU is allocated a /12 in the IPv6 address space to act as an alternative IP address registry, as other reports have suggested.[12]

TABLs could be implemented by ITU on a global basis, or could be implemented based on the country approach. A global approach to TABL implementation by ITU has

[12] Ramadass (2009).

many advantages. If ITU offered TABLs on a uniform, global basis, its ability to offer an attractive alternative to the regionalized, needs-based policies and practices of the RIRs would be greatly enhanced. Global demand for TABLs is likely to be thicker and more liquid than if it were confined to a single region or country.

Another option, however, is to treat transferable leases as a way of re-allocating address blocks assigned to Country Internet Registries (CIRs). If ITU member-states create a regime of CIRs, then it is possible that the initial allocation method might give some countries far more than they need and other countries less. So in addition to dividing up its portion of the IPv6 address space into 200+ country blocks to be allocated to CIRs, the ITU could reserve a sizable chunk, apply the TABL principle to it, and allow CIRs that thought they needed more addresses than the initial ITU allocation to lease additional blocks from that pool. In this case, the right to acquire, sell and transfer TABLs could be restricted to CIRs.

5. Appendix 1: Outline of the evolution of IPv4 institutions and policies

"Address space exhaustion is one of the most serious and immediate problems that the Internet faces today." (RFC 1335, May 1992)

Although this report is concerned with the IPv6 routing-addressing space, IPv6 enters a world of institutions and policies formed around IPv4. Current approaches to management of the routing-addressing space can best be understood through a historical analysis of the way they evolved in response to specific problems at specific moments in time. IPv4 address allocation has gone through four stages, which can be identified as follows: 1) Legacy classful allocations; 2) The replacement of class-based allocations with Classless Inter-Domain Routing (CIDR) and the emergence of regional Internet registries (RIRs); 3) institutionalization of the RIRs and formalization and tightening of needs-based assessments and fees; and 4) the depletion of the IPv4 address pool. For ease of reference, the outline follows RFCs, citing the critical documents.

5.1. Legacy classful allocations (1981 – 1990)

a. RFC 791 (September 1981) defines IPv4 address space.
b. Address space divided into three classes, A (16.7 million addresses), B (16.7 thousand addresses) and C (256 addresses).
c. Internet registry run by the Defense Data Network's Network Information Center (DDN-NIC); address registration function was contracted out to Stanford Research Institute

d. IANA function established at USC Information Sciences Institute by Jon Postel.
e. Address allocations made largely upon request; scarcity and the possible depletion of the address pool was not a major consideration.
f. No contractual agreement with assignees and no plans for recovery or reclamation of unused resources.

5.2. First response to scaling problems (1991 – ~1997)

a. First Regional Internet Registries created (RIPE-NCC, APNIC). Rationale set out in RFC 1174 (August 1990).
b. Scaling problems arise in routing table growth, and address space depletes. RFC 1466 (May 1992).
c. Classless Inter-Domain Routing (CIDR) implemented. RFC 1519 (September 1993)
d. Provider-based leasing model approach to routing. RFC 2008 (October 1996)
e. Role of RIRs formalized and policies and practices incorporate CIDR and provider-based aggregation. RFC 2050 (November 1996).

5.3. Maturation of the RIR regime (~1998 – 2007)

a. Full privatization of the central IR (ARIN) and partial privatization of IANA (ICANN)
b. Formal contractual regime and fee structures put into place
c. RIRs become stable, well-funded organizations
d. Formation of LACNIC (October 2002) and AFRINIC (April 2005)
e. Continuity in policy, but tighter controls

5.4. Response to impending exhaustion of free pool scarcity phase (2008 – ongoing)

a. Transfer markets

b. Reclamation efforts

c. Inter-regional transfers via IANA

d. Reserve blocks for various purposes

e. Encouragement of migration to IPv6

6. Appendix 2: Aggregation fees

We considered, but ultimately do not recommend, modifying the proposed system of address block fees to encourage route aggregation. Earlier discussions of the routing tragedy of the commons explored the notion of charging for route announcements. (Rekhter, Resnick and Bellovin, 1996) These explorations led many to conclude that it was impossible to do so, because it was assumed that ISPs would have to directly negotiate and compensate each other for route announcements on a bilateral basis.

It is, however, possible to reward route aggregation (or conversely, to penalize de-aggregation) by linking the address block fees to aggregation efficiency. The fee for initial allocations would be based strictly on the size of the block, and the periodic recurring fee would be based on both the size of the block and on the address block holder's aggregation efficiency. Recipients of address allocations could pay recurring charges on an annual or quarterly basis.[13] Recurring charges would be based on the size of the block with an added charge for *the number of routes* an organization announces per allocated block. E.g., if only 1 route is announced into the exterior routing domain for the allocated block, perhaps the recurring charge is 0; then it increases until it reaches its highest level when the routes announced are completely disaggregated to the /64 level. Such a fee would have to vary depending on the size of the block, increasing for smaller blocks and decreasing for larger ones. (This aspect requires further exploration.) Such a fee structure might combine and integrate the incentive to conserve with the incentive to aggregate routes.

---

[13] From this high-level sketch, the period does not matter much; from a practical administrative standpoint it might matter very much.

The purpose of such a fee structure would be to internalize the externality associated with routing table growth. Such an approach "taxes" deaggregation. Calibrating recurring fees to reward aggregation efficiency might make it possible to have a more liberal trading regime that would allow holders of TABLs to sell portions of their blocks, because it would ensure that trading of deaggregated blocks would occur mainly when it *increased* routing efficiency (or at worse, left it the same). Address block trades that led to deaggregation would increase the trading parties' recurring costs. The costs of deaggregation would overpower initial allocation costs unless extreme scarcity developed in the ipv6 space, which seems highly unlikely for the next 50-60 years at least.

We do not recommend an aggregation fee, for two reasons. First, it is unclear whether the problem of routing table bloat is serious enough to justify it. Currently, ISPs can filter prefixes that exceed a certain length if they are concerned about that problem, and as noted earlier the projected growth of the routing table may fall within the technological capacity of improvements in router capacity. Second, the number of routing announcements often reflects traffic engineering concerns. That is, organizations and ISPs use route announcements to steer traffic over specific link facilities. Traffic engineering can be as important to the efficiency of the Internet as route aggregation. Unless a serious crisis of routing table scalability emerges, it is probably better to avoid any route aggregation fees and allow ISPs to make these tradeoffs on their own.

7. References

Cowie, J. (2009). *The adventurous parts of the internet.* Retrieved July 1, 2009, from http://www.renesys.com/blog/2009/01/the-adventurous-parts-of-the-i.shtml

Euopean Union. *IPv6: Enabling the information society*. (2009). Retrieved July 15, 2009, from http://ec.europa.eu/information_society/policy/ipv6/index_en.htm

Fuller, V., Li, T., Yu, J., & Varadhan, K. (September 1993). *Classless inter-domain routing (CIDR): An address assignment and aggregation strategy*. RFC 1519, Internet Engineering Task Force.

Gerich, E. (May 1993). *Guidelines for management of IP address space.* RFC 1466, Internet Engineering Task Force.

Hinden, R., & Haberman, B. (October 2005). *Unique local IPv6 unicast addresses*. RFC 4193, Internet Engineering Task Force.

Hubbard, K., Kosters, M., Conrad, D., Karrenberg, D., & Postel, J. (November 1996). *Internet registry IP allocation guidelines.* RFC 2050, Internet Engineering Task Force.

Huitema, C. (November 1994). *The H ratio for address assignment efficiency*. RFC 1715, Internet Engineering Task Force.

Huston, G. (2009). *BGP in 2008.* Retrieved June 15, 2009, from http://www.potaroo.net/ispcol/2009-03/bgp2008.html

Huston, G. (2005). *Just how big is IPv6? Or, where did all those addresses go?* Retrieved June 20, 2009, from http://www.potaroo.net/ispcol/2005-07/ipv6size.html

Huston, G. (2001). Analyzing the Internet BGP routing table. *The Internet Protocol Journal*, *4*(1), 2-15.

IAB & IESG. (September 2001). *IAB/IESG recommendations on IPv6 address allocations to sites.* RFC 3177, Internet Engineering Task Force.

Lehr, W., Vest, T., & Lear, E. (2008). Running on empty: The challenge of managing internet addresses. Paper presented at the 36th Annual Telecommunications Policy Research Conference, September 27, 2008. Arlington, Virginia, USA.

Meng, X., Xu, Z., Zhang, B., Huston, G., Lu, S., & Zhang, L. (2005). IPv4 address allocation and the BGP routing table evolution. *ACM SIGCOMM Computer Communication Review*, *35*(1), 71-80.

Meyer, D., Zhang, L., & Fall, K. (September 2007). *Report from the IAB workshop on routing and addressing.* RFC 4984, Internet Engineering Task Force.

Millet, S., & Huston, G. (August 2005). *Proposal to amend APNIC IPv6 assignment and utilisation requirement policy*. Proposal No. prop-031-v001.

Mueller, M. (2008). *Scarcity in IP addresses: IPv4 address transfer markets and the regional internet address registries.* Internet Governance Project. Retrieved July 15, 2009, from http://internetgovernance.org/pdf/IPAddress_TransferMarkets.pdf

Mueller, M., & Kuerbis, B. (2008). *Regional address registries, governance and Internet freedom.* Internet Governance Project. Retrieved July 15, 2009, from http://internetgovernance.org/pdf/RIRs-IGP-hyderabad.pdf

Narten, T. (June 2005). *Issues related to the management of IPv6 address space.* IETF Internet Draft. Retrieved from http://ietfreport.isoc.org/all-ids/draft-narten-iana-rir-ipv6-considerations-00.txt

Ostrom, E. (2005). *Understanding institutional diversity*. Princeton, NJ: Princeton University Press.

Postel, J. (September 1981). *Internet protocol.* RFC 791, Internet Engineering Task Force.

Ramadass, Sures. (2009). A Study on the IPv6 Address Allocation and Distribution Methods. National Advanced Center for IPv6 Center of Excellence, Malaysia.

Rekhter, Y., Resnick, P., & Bellovin, S. (1996). Financial incentives for route aggregation and efficient utilization in the internet. In Kahin and Keller (ed.), *Coordinating the Internet*, Cambridge, MA: MIT Press.

Velde, G. V. d., Popoviciu, C., Chown, T., Bonness, O., & Hahn, C. (December 2008). *IPv6 unicast address assignment considerations.* RFC 5375, Internet Engineering Task Force.