A Primer on IPv4 Scarcity

Philipp Richter
TU Berlin / ICSI
prichter@icsi.berkeley.edu

Mark Allman
ICSI
mallman@icir.org

Randy Bush
Internet Initiative Japan
randy@psg.com

Vern Paxson
UC Berkeley / ICSI
vern@cs.berkeley.edu

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

arXiv:1411.2649v3 [cs.NI] 27 Feb 2015

ABSTRACT

With the ongoing exhaustion of free address pools at the registries serving the global demand for IPv4 address space, scarcity has become reality. Networks in need of address space can no longer get more address allocations from their respective registries.

In this work we frame the fundamentals of the IPv4 address exhaustion phenomena and connected issues. We elaborate on how the current ecosystem of IPv4 address space has evolved since the standardization of IPv4, leading to the rather complex and opaque scenario we face today. We outline the evolution in address space management as well as address space use patterns, identifying key factors of the scarcity issues. We characterize the possible solution space to overcome these issues and open the perspective of address blocks as virtual resources, which involves issues such as differentiation between address blocks, the need for resource certification, and issues arising when transferring address space between networks.

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network Operations—*Network Management*; C.2.2 [**Computer-Communication Networks**]: Network Protocols—*Protocol architecture (OSI model)*

Keywords

IPv4 address exhaustion; IPv6 transition.

1. INTRODUCTION

The Internet's design philosophy has facilitated enormous, rapid, and de-centralized growth, from a specialized research facility to a massive network of global importance. In turn, the tremendous growth enabled by the original design also outpaced engineers, researchers, and policy makers. This is clear in the numerous technical challenges that have arisen—from the lack of support for traffic engineering and routing security, to the scalability issues of the initial mapping of hostnames to IP addresses, to the lack of congestion control, to the inability to accommodate mobility.

The community has largely been able to address such issues, albeit not always in the most elegant way given that changing a running system presents a challenging target in many cases. However, under the surface, policy and governance issues arose. While scientists and engineers often ignore these issues, they ultimately shape what we can deploy in production. In this paper we consider the entanglement of policies and governance with the technology for one of the Internet's key resources: IP addresses.

Transmission of data between hosts across the Internet requires network layer addresses to name the endpoints—i.e., IP addresses. In IP version 4, an address is represented by 32 bits in the IPv4 header; hence there is a finite pool of roughly 4B addresses available. The network routing system cannot keep enough state to deal with each individual address and therefore aggregates addresses into blocks. Originally blocks were allocated quite informally, but as the Internet grew the complexity of the process did as well. At this point the address space is nearly entirely allocated.

The Internet engineering community has long recognized the impending exhaustion of IPv4, and in response designed a replacement network-layer protocol with much longer addresses, IPv6. However, given the network layer's critical functionality, deploying IPv6 has proven difficult.[1] Therefore, it is now widely acknowledged that IPv4 will continue to play a significant role for a long time within the confines of the current resource limits. While there are technical maneuvers we can still make to cope—e.g., adding layers of network address translation (NATting)—IPv4 address blocks have already become a good that people exchange on secondary markets. This reality brings yet another challenge to the Internet's policy and governance ecosystem.

In this paper we first survey the evolution of the community's management of IP addresses, for which we include both a discussion of the relevant policy structures and organizations, as well as empirical illustrations of the allocation and use of IP addresses over time. This history then leads to a set of observations about protocol design and the accompanying stewardship of community resources.

2. EVOLUTION OF ADDRESS MANAGEMENT

From its standardization in 1981 [69] until now, the management of IP addresses has undergone drastic change. The changes were mainly a result of the evolution of the Internet from a research network to a global commercial network and the corresponding need to establish international frameworks to manage its critical resources. We elaborate on this evolution in three time phases: The *Early Registration* Phase starting with the arrival of IPv4, the *Needs-based Provision* Phase leading to the modern registry framework, and the recently entered *Depletion and Exhaustion* Phase.

[1] See [32] for an understanding of IPv6 deployment from multiple viewpoints.

Juniper Exhibit 1006, page 1

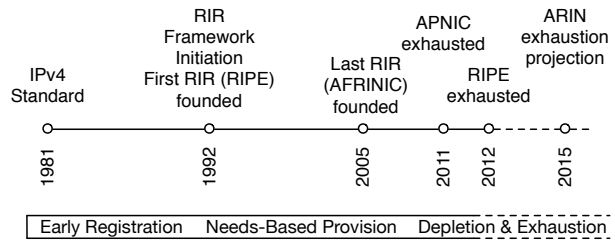

Figure 1: Evolution of address management.

2.1 First Phase: Early Registration.

Initially, address blocks were allocated quite informally, with Jon Postel serving as the "czar" personally attending to each allocation. Postel periodically re-published RFCs enumerating the current address assignments ("please contact Jon to receive a number assignment") [68]. At that time, addresses block allocations came in one of three *classes*: class A networks (2^{24} addresses), class B (2^{16}), and class C (2^8). Classful addressing required a *network identifier* of one of these distinct types, meaning that an operator requesting significantly more addresses than provided by a particular threshold would instead be allocated a larger class network. Given the coarse-grained nature of the differences between these classes, this policy led to heavy internal fragmentation and thus waste of address space.

Early (1981) in the Internet's evolution, parties had already registered 43 class A networks, allocating in total more than 700M addresses [68]—vastly larger than the number of hosts actually connected at that time.[2] While scarcity in address blocks was not mentioned as a looming issue, the notion of different sizes of networks (A, B and C) suggests early recognition of the finite nature of network address blocks and the need for some sort of stewardship when parceling them out to different parties. The responsibility for the management of address space led to formalizing the notion of the IANA (first mentioned in IETF documents in 1990 [72]), and, in the same timeframe Solensky, drawing upon allocation statistics, predicted IPv4 address exhaustion in the late '90s [86].

2.2 Second Phase: Needs-based Provision.

The need for a more distributed and parsimonious framework to allocate IP addresses—shaping the modern registry structure—appeared at least as early as 1990 [31], with further refinements in 1992 and 1993 [41]. The discussion at that time included the need to distribute the administration of IP address blocks to regional registries, covering distinct geographic regions to better serve the respective local community—consciously fragmenting the registry. In addition, classless inter-domain routing (CIDR)[3] and private address space[4] arose in 1993–4 to further conserve publicly routable address space.

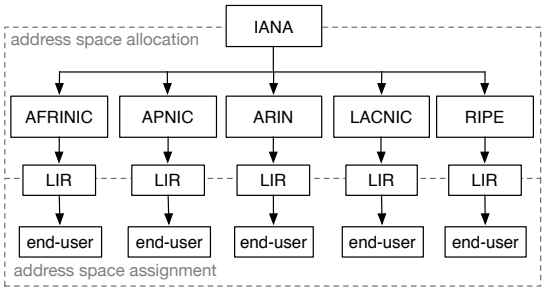

Figure 2: Regional Internet Registry system.

The modern framework of Regional Internet Registries (RIRs), established in the years between 1992 and 2005, was very specific that conservation of address space was a primary goal [43]. Five RIRs emerged, run as non-profit organizations: RIPE for Europe in 1992, APNIC for the Asia-Pacific in 1993, ARIN for the North-Americans in 1997, LACNIC for Latin America in 2002, and AfriNIC for Africa in 2005.

The RIRs manage the distribution of IP address resources, each according to their local policies. Policies within the RIRs are created using a community process; for details of the process for each RIR, see [5, 9, 13, 53, 77]. For the most part, anyone can submit an RIR policy proposal which then undergoes an open discussion and review process, usually carried out on mailing lists as well as in working group and policy meetings. Adopting a proposal requires the community to reach a degree of consensus as reflected in these discussions.

We sketch the structure of the RIR framework in Figure 2. The IANA serves as the parent organization, *allocating* large free address blocks (/8, i.e. 2^{24} addresses, granularity) to an RIR once their regional free pool reaches a low threshold level. The RIRs then further allocate subsets of these address blocks to their members, the so-called LIRs (Local Internet Registries), which are mainly ISPs. The LIRs then *assign* address blocks to either smaller ISPs or for their own infrastructure. Thus, the allocation of a block reserves it for (future) use, while the assignment of parts of an allocation puts that subset into use.[5] ISPs decide for themselves whether to become LIRs—meaning entering a direct contractual relationship with the respective RIR—or to rely upon their upstream provider to assign address space to them.[6]

[2] Address registration statistics in terms of number of blocks and block holders varied heavily among the first published RFCs.

[3] CIDR [39] supported routing and forwarding on bit-aligned, as opposed to the previous byte-aligned, variable-length prefixes. CIDR denotes prefixes as a combination of an IP address and a corresponding network mask, such as *1.1.2.0/23* specifying a network with 2^9 IP addresses that share their top 23 bits. Introducing CIDR required significant network restructuring efforts, as well as changes to routing protocols and hardware (see, for example, [38]).

[4] Reserved address blocks not globally routable, and thus usable concurrently within multiple networks as long as the given hosts do not require globally reachable IP addresses [71].

[5] The APNIC and LACNIC regions also have *National Internet Registries* (NIRs), which act as intermediaries between the RIR and the LIR to serve specific countries. For example, JPNIC does so for Japan.

[6] Under some circumstances, RIRs can also assign address space directly to end users—so-called provider independent (PI) address space. Such assignments usually arise due to the user's need to connect to multiple upstream providers (multi-homing), and thus requiring independent address space. For more details, see for example § 4.2 in the ARIN NRPM [14] or the RIPE policy documents [83]. For a practical guide for operators, see [27].

Juniper Exhibit 1006, page 2

During the needs-based provision phase, one of the key principles was that receivers of address space (LIRs) must *justify* their need for the address blocks they receive, though some RIRs no longer require this in some contexts (e.g., RIPE for "last /8" allocations—see below). LIRs requesting new allocations had to provide documentation showing a sufficient *utilization rate* of prior allocations, namely that a given proportion of prior allocations were assigned to end-users, as well as documentation of the intended use of new allocations. RIRs might also request more detailed information, such as how many and what type of hosts were connected to assigned subnets. LIRs passed these policies on to their end-customers. For example, if a customer of a transit provider required blocks of IP addresses, they had to fill out corresponding LIR-specific forms detailing the intended use of that block (e.g., [66]).

The global nature of the Internet raises the question of when an organization is supposed to be served by a specific geographic region. Whether or not a company can become an LIR under a specific RIR is not explicitly stated, but is usually determined by the registered address of a company. However, there also are organizations with multiple subsidiaries as members of—and holding address resources from—multiple RIRs [45]. While address blocks are theoretically assigned and "used" by organizations operating inside the region of the allocating RIR, current policies are inconsistent regarding explicit constraints on the geographic region of an address block's actual use in the sense of where connected devices reside.[7]

2.3 Third Phase: Depletion and Exhaustion.

The five RIR communities agreed to a policy regarding address block allocation upon the onset of exhaustion, which ICANN—the international body responsible for the IANA function—ratified in 2009 [44]. The policy dictated that when the IANA's IPv4 free pool reached five remaining /8 blocks, the IANA would distribute these blocks simultaneously and equally to the five RIRs. In February 2011, the IANA allocated its last five free /8 address in accordance with the policy, one to each RIR [65]. After that point, from a global perspective the pool of available IPv4 addresses was fully depleted.

Once the RIRs started to allocate from this last block from the IANA, the "last /8" policies introduced by each RIR went into effect (e.g., APNIC's per [6]), imposing more restrictive allocation policies to further conserve this final address block and to allow new market entrants to still receive a last allocation, e.g., to implement IPv4-to-IPv6 transition mechanisms. Thus, LIRs could receive a single (small) allocation from this block. This transition occurred in April 2011 for APNIC, in September 2012 for RIPE, and in June 2014 for LACNIC, upon the exhaustion of their respective free pools. ARIN's exhaustion date is likely to occur in early 2015, while AFRINIC's pool should last until 2019 [40].[8] LIRs in need of address space now need to find other means of obtaining address space.

[7] ARIN has a policy proposal to explicitly allow out-of-region use [15], and a RIPE official stated that RIPE permits out-of-region use, assuming that the address blocks originate at some point from within the RIPE region (e.g., by a router at a European Internet Exchange Point) [74]. Numbering resources under the stewardship of LACNIC must be distributed among organizations legally constituted within its service region, and mainly serving networks and services operating in this region. The AFRINIC community, on the other hand, has discussed explicitly limiting out-of-region use to prevent possible exploitation of their IP address resources from operators in other regions [4].

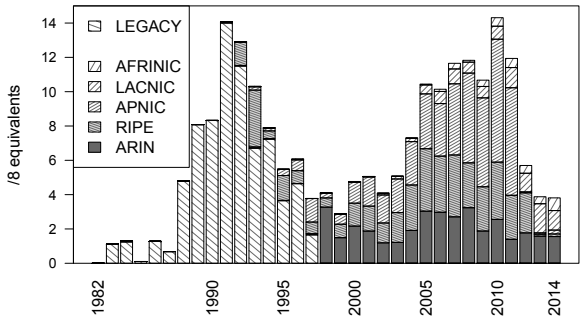

Figure 3: Yearly allocations of IPv4 address blocks.

3. EVOLUTION OF ADDRESS BLOCK ALLOCATION

Per the above, almost all of the free IP address blocks have been distributed. We can group today's address blocks into three categories: (i) blocks given out prior to the RIRs' existence, termed *legacy* address space;[9] (ii) blocks given out during the era of the RIRs, termed *allocated* address blocks; and (iii) *reserved* address blocks, such as those set aside for multicast and private addressing.

Figure 1 on page 2 shows a timeline of the most significant events in the evolution of address block allocation. One cannot pinpoint the transition between the above-mentioned phases precisely: the RIRs were founded years apart, hence ISPs in some regions received legacy address space for a longer period than in other regions. ARIN, for example, began in 1997, whereas RIPE was founded in 1992. Thus, address space holders in the European region received allocated address blocks earlier while holders in North America were still receiving legacy blocks. The transition between phase 2 and 3 is ongoing, as two RIRs (ARIN and AFRINIC) still have unexhausted free pools.

In the remainder of this section we present an empirical lay-of-the-land of the state of these allocations.

3.1 History of Address Block Allocations

The IPv4 address space consists of 2^{32} possible addresses, an equivalent of 256 /8 address blocks. Of these 256 /8 blocks, 35.3 are reserved by the IETF, e.g., for multicast, private use, and future use. This leaves 220.7 /8s worth routable address space.

In the following, we present a historical view on IPv4 address consumption from an RIR allocation point-of-view. We rely on allocation files provided by the RIRs [64]. Figure 3 shows the address blocks given out by the registries over the years, as well as those given out prior to the existence of the modern RIR framework (shown as LEGACY).

Two peaks in address consumption are quite visible: The

[8] We set the exhaustion date to when the RIRs started to allocate from their last /8, consistent with [40].

[9] LACNIC (and possibly AFRINIC) uses the date of ARIN's inception as their "legacy" threshold, not their own formation, as they would otherwise be unable to apply their policies to addresses that predate their formation.

Juniper Exhibit 1006, page 3

	handed out /8s	of which legacy /8s	available /8s
ARIN	100.5	∼ 64.9	0.35
RIPE	47.6	∼ 11.93	0.97
APNIC	51.0	∼ 4.40	0.74
LACNIC	10.9	∼ 0.58	0.20
AFRINIC	4.5	∼ 0.02	2.63
total	**214.5**	**∼ 81.83**	**4.89**
% of routable	**97.2%**	**∼ 37.1%**	**2.2%**

Table 1: Address space statistics (February 2015).

first occurs in the "Early Registration Phase" in the late 80's and early 90's. As discussed in the previous section, address space conservation was not yet a primary concern, and classful allocations resulted in heavy internal fragmentation of address space. The allocation rate drastically decreased in subsequent years, as address space conservation was implemented by the RIRs. Address consumption rates in the late '90s and early 2000s suggested IPv4 address exhaustion would not happen before 2020. The second peak, starting in the mid-2000s was dominated by allocations in the APNIC region, and comprised more than 50% of all allocations in 2010 and 2011. After the exhaustion of the IANA free pool in 2011, a rapid decline in further allocations in 2012 is quite visible. Currently, fewer than 6 /8 equivalents are available for distribution by the RIRs.

The responsibility for the administration of legacy address blocks was transferred to ARIN upon its inception in 1997 [60]. ARIN subsequently re-distributed some of these legacy blocks to the various other RIRs for respective holders located outside the ARIN region. This happened in the course of the ERX (Early Registrations Transfer) project [73]. Yet, most legacy address space is still administered by ARIN, a symptom of North America's dominance of the early Internet.

Table 1 gives an overview of the distribution of the address space among the RIRs (in February 2015). The first column is the number of /8 equivalents, as listed in the allocation files of the RIRs. The second column is an estimate of how much address space is legacy (given out in Phase 1) for each RIR.[10] The last column shows the number of /8s per RIR that are available for allocation. We observe that close to 97% of the IPv4 address space has already been allocated, with less than 3% available for further allocation. Some address blocks are in a reserved state (e.g., for temporal assignments for Internet experiments or conferences), and thus neither available nor handed-out. The heavy allocation rates in the last years prior to exhaustion mainly reflect heavy consumption in the APNIC region. This could reflect a degree of hoarding, but might simply reflect booming Internet deployment in Asia.

3.2 History of Routing

In the last section we outlined how the management of IP addresses evolved over time. From a pure allocation perspective, the address space is now close to fully exhausted. One important question is the degree to which allocation reflects actual use. We can consider this in two parts: (1) the degree to which elements of allocated blocks are routed, and thus potentially in use; (2) the degree to which addresses within routed blocks are in fact used. Here we assess the first of these, as we can much more readily obtain insight into it (via the global routing table as publicly available from the RouteViews project) than we can for the second consideration.

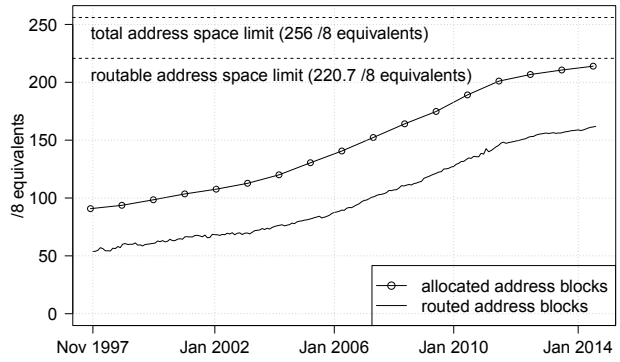

Figure 4: Allocated and routed address blocks.

Figure 4 shows the number of routed address blocks (expressed as /8 equivalents) over the last 16 years, along with the cumulative total of allocations made by the RIRs. We see that by 1997 more than 25% of the routable address space was advertised, which gradually increased to over 70% in January 2014. While there is an increasing trend in the '00s, in the last two years the rate has been fairly stagnant, perhaps reflecting address exhaustion. It should be noted that the growth of the Internet in its early prime, starting in 1997, used some 50% of the available address space, while the 25% routed prior to that time is likely due to classful allocations and rather lax allocation policies.

Figure 5 shows the evolution of the routed address space from 1997 until 2014, by plotting for each /8 the fraction of routed addresses, ranging from white (no address blocks advertised) to black (all address blocks advertised), with the various ranges annotated according to their address types.[11]

The most striking observation from this plot is that the use of address blocks is very unevenly distributed. Address ranges assigned prior to the existence of the RIRs, the legacy ranges, exhibit much fewer routed address blocks, whereas the RIR-allocated ranges show a gradually increasing and consistent routing pattern. We see that the measures taken in Phase 2, namely the delegation of finer-grained address blocks (CIDR), together with the address conservation principles of the RIRs, indeed had noticeable effect. Hence, efficient address management greatly improved the utilization of address space, but did not enhance utilization in legacy ranges outside of their scope of operation. Today, address blocks in the legacy range have the greatest supply of free and usable address space. In fact, as of February 2015 more than 90% of the allocated address space is routed but only some 50% of legacy address space.

[10] For ARIN, we consider all address blocks handed out prior to December 1997 as legacy. For the other RIRs, we consider all address blocks transferred as part of the ERX project as legacy, in addition to blocks *25/8*, *51/8*, *53/8* and *57/8* for RIPE and *43/8* for APNIC. Some of these blocks may have been voluntarily returned or otherwise changed their status. Thus, the number of legacy blocks only serves to give a sense of the landscape.

[11] A few /8 *legacy* block ranges of former class A networks were not given out, and are thus allocated. In addition, some smaller address blocks in the former class B range were allocated by the RIRs, hence the notation "mainly" in the figure.

Juniper Exhibit 1006, page 4

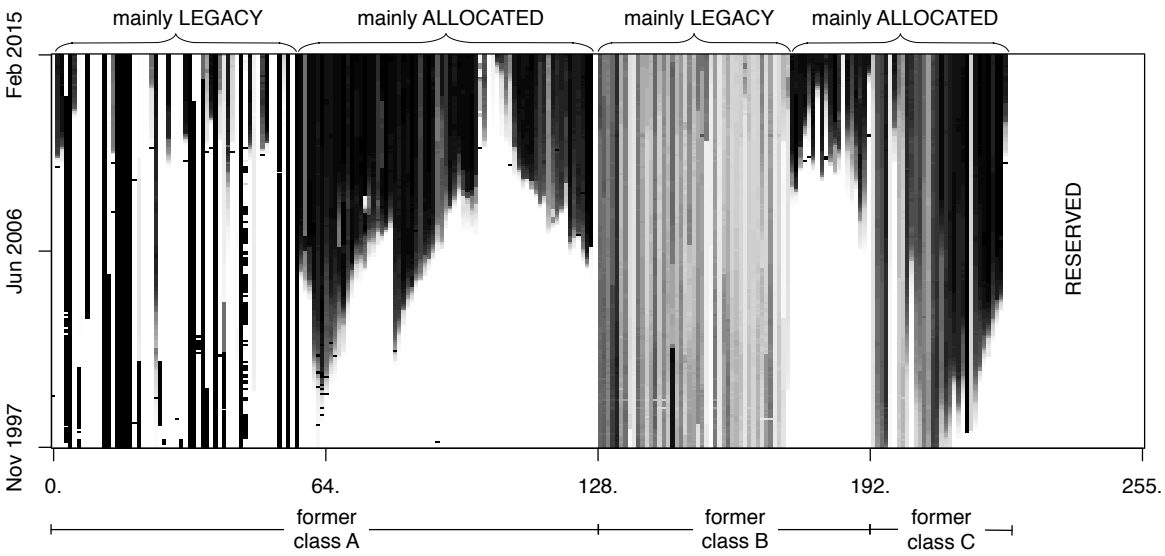

Figure 5: Evolution of the distribution of routed address blocks over the total address space.

The caveat when using routing tables to reason about the utilization of address blocks is that, while it gives an indication of address space use (clearly visible here), a routed address block does not necessarily mean that it is in active use. Recent estimates range from 47% to 60% [33, 34, 89] of routed /24 address blocks that are actually used, meaning that they are actively engaged in communication. Actual use of address blocks can be measured actively (e.g., probing every IP address with a *ping*) or by relying on passive measurements (e.g., identifying those parts of the address space that actively engage in communication—emitting traffic). Zander et al. [89] also used meta-information such as Wikipedia edit logs and applied a statistical model in order to account for address blocks that are not detected by such methods. While an address block being routed does not imply its actual use, unrouted address blocks, on the other hand, might be in private use for interconnecting networks not publicly reachable.

Hence, while the IP address space is close to fully exhausted, from an allocation perspective, scarcity seems to be less of an issue from a purely technical perspective (e.g., routing). While it requires further work to quantify "efficient use", we can clearly see significant differences between legacy address space and allocated address space.

4. IP ADDRESSES AS A RESOURCE

IP addresses are virtual resources. In this section, we elaborate associated issues.

4.1 Addresses: All The Same, Only Different

At first, one might consider IP addresses as a fully homogeneous (fungible) resource, but in fact not all addresses have equivalent properties. First, the size of a given address block governs its routability. Larger address blocks are less likely to be filtered by other operators, and can be deaggregated into smaller entities, allowing networks to better engineer their route announcements. In addition, an address block comes with *history*: for example, a block previously used by spammers will more likely be found on blacklists, limiting one dimension of its usability. Finally, the properties of address blocks differ depending on their allocation standing and any associated policy restrictions, as noted in the next section.

4.1.1 The case of allocated address blocks

Allocated address blocks given out by the RIRs (Phase 2) are contractually constrained—in a more or less explicit way—as not constituting the *property* of the respective holder. ARIN, RIPE and AFRINIC have explicit "no property" statements in the documents a receiver of address space must agree to [2, 22, 78], while LACNIC and APNIC have more implicit statements in their contracts, not mentioning ownership or property by name. LACNIC states that it can withdraw address blocks from holders [55] and APNIC states that it [only] hands out resources on a "license basis" [10]. The RIRs apply different policies for address space they give out, both with regard to the requirement to document how address space is used as well as with regard to transferability of address blocks. Hence, for RIR-allocated address blocks, the holder will generally have to agree to policies and eventual policy changes as imposed by the respective RIR. Thus, the region associated with an address block directly affects the policies that govern it and thus also its value.

With respect to the possibility of RIRs unilaterally *reclaiming* unused address space from LIRs, the policy documents differ. ARIN clearly rules out unilateral reclamation in its current RSA [22]. APNIC does not mention this possibility by name in its documents, but states that *"If an allocation or assignment becomes invalid then the address space must be returned to the appropriate IR"* [10]. AFRINIC states the possibility of *"revocation or withholding of the service supplied"* [2], and RIPE that it might deregister resources if members fail to comply with their policies [78]. We are not aware of any cases of a unilateral reclamation of

Juniper Exhibit 1006, page 5

allocated address space to date, aside from those where an address holder went defunct without successor.

4.1.2 The case of legacy address blocks

Legacy address blocks, on the other hand, are not in general governed by contractual requirements imposed by any RIR. A noteworthy point with regard to IP addresses as resources is the ongoing discussion whether IP addresses can be considered property or not [84]. Per Figure 5, much of today's unrouted address space is legacy, and thus not considered to be subject to current RIR policy.

The RIRs do maintain the registry databases and the anchors for reverse DNS mappings for legacy blocks. However, the attitude of the RIRs towards holders of legacy resources varies. In the course of the last decades, the RIRs—mainly ARIN [20]—started several initiatives to contact holders of legacy address space with the goal of establishing some contractual agreements between the holder and the RIR. As the documentation of legacy allocations is often poor (e.g., outdated information), many holders of legacy resources might not even be approachable. ARIN offers LRSAs (Legacy Registration Services Agreement) [21] to holders of legacy address space in their region. LRSAs establish a more formal relationship between the address holder and ARIN, contain an explicit "no property" clause, and also contractually obligate the legacy holder to ARIN's policies, including the policy for transfer to other entities (or when the holder requests additional address space from ARIN). In late 2007 ARIN sent out more than 18K letters to legacy holders [20]. Their data shows that as of 3 years later, fewer than 1,000 LRSAs were in turn requested by the holders, and LRSAs cover less than 15% of the legacy address space in the ARIN region [18]. One address broker publicly suggests to legacy holders to not sign such LRSAs [50]. Another ARIN document states "*All of the IP address space that ARIN administers, including legacy space, is subject to ARIN policy*" [25]. RIPE, on the other hand, adopted a proposal in February 2014 to offer registration services to holders of legacy address space and not impose particular regulations on transfers of registered legacy address blocks [80].

Regarding the possibility of *reclaiming* unused legacy address blocks, ARIN states that it will not attempt to unilaterally reclaim legacy address space [19]. APNIC and RIPE ran initiatives to contact holders of *legacy* address blocks to recover address space [11, 81] but left the decision up to the respective holder. In case of the RIPE initiative, 400 holders were contacted of which 16 returned address space to RIPE. However, there are prominent examples of voluntarily returned legacy address blocks, such as Stanford University voluntarily returning its /8 legacy address block in 2000 [62], as well as some other organizations [46].

A meeting convened by ICANN in 2012 informally addressed issues related to legacy address resources [45]. The discussion involved representatives from the RIRs, network operators holding legacy and non-legacy address resources, and address brokers. On one hand, it was argued that legacy resources by their nature do not differ from other IP address blocks, and should thus be subject to the same policies. On the other hand, holders of legacy address space argued that *grandfathering* applies—meaning that as legacy address space was given out prior to RIR policies, they are not subject to any policies subsequently created by RIRs.

Hence, the open question with regard to legacy holders is whether they are bound to the terms of the registry that currently provides registration services to them—in a more general way, whether they hold ownership rights for their addresses or not.

4.2 Resource Certification & Enforcement

In the case of IP addresses, no global system exists to either authoritatively verify the ownership of a given address block nor to prevent the usurping of address blocks by illegitimate users. Inter-domain routing as instantiated by BGP does not itself provide any mechanisms to ensure routing only by a block's legitimate holder. While the community readily recognizes BGP's lack of security features, including its inability to authenticate routes, a large body of research and accompanying deployment efforts has done little to change this situation in productive environments (see [30] and references therein).

The RIRs publicize the mapping of address spaces to their respective holders via registry databases (WHOIS), which can be queried publicly, and by delegating the respective reverse-DNS zones (`.in-addr.arpa`) to authoritative nameservers specified by the address holders. This latter enables the holders to specify PTR records for IP addresses in the respective namespace (not a fundamental requirement or hallmark of ownership, but certainly operationally useful). Neither of these mechanisms provide sufficient information to directly validate (or invalidate) route advertisements, such as by authoritatively indicating the origin AS.[12] Thus, the administrative management of address space is largely decoupled from its actual use.

The degree to which a prefix is usable by some entity—and which entities have the capability to use it—simply depends on how far a route advertisement for the given prefix propagates, which directly translates into how many hosts on the Internet can interact with hosts in the given address block.

The propagation or non-propagation of prefix advertisements depends on the route filtering performed by the border routers of ISPs. To configure these filter settings, the community has established routing registries (IRR), where network operators can register route objects to express prefix ownership in the form of prefix-AS mappings [30]. The various IRR databases are managed by several independent organizations, including ISPs, RIRs and others [51]. However, not all address space is registered in some registry (only around 50% according to [87]) and information in these registries is known to be significantly inaccurate [52]. Many IRRs allow their participants to introduce essentially any route object without further validation [88]. Complications with the IRR can again result in ISPs not filtering advertisements from their peers using IRR information at all [36].

There are well-known cases of erroneous IP address block advertisements, be it hijacking of address blocks by spammers [70] or advertisements caused by misconfigurations. As an example, a Pakistani ISP erroneously advertised a prefix belonging to YouTube in 2008, resulting in an extensive global outage for that service [61].

The Internet Engineering Task Force (IETF) has developed a solution to this problem based on the RPKI (Resource Public Key Infrastructure) [56]. The basic function of the RPKI is to provide cryptographically verifiable attes-

[12] The ARIN WHOIS database recently started to provide a field for the origin AS, but the field is often unset and prominent cases of inconsistencies exist [67].

Juniper Exhibit 1006, page 6

tations to address space and AS number allocations using a X.509 based hierarchy. RPKI uses the IANA and the RIRs as trust anchors, which give out certificates for resources they manage. Currently, RPKI services are offered by the RIRs as a free opt-in service only to their members.[13]

Based on the RPKI database, routers can verify that an AS advertising a specific prefix is in fact authorized to do so, which is referred to as RPKI-based origin validation [28]. This only prevents accidental advertisements and is not intended to prevent malicious attacks, as the full AS path is not validated but only the origin of the path [29].[14] RPKI is supported by current routers from Cisco, Juniper, and Alcatel-Lucent, yet currently only around 5% of the routable address space is covered by RPKI and by far the largest share of that address space is in the RIPE region [79].

While the problem of securing the advertisement of a prefix by only the respective holder is well-known and many approaches have been proposed over the years, little has changed in productive environments. Faced with the increasing scarcity of IP addresses (and the corresponding increasing value of addresses as resources), a functional scheme for certifying resources will be a key requirement in the near future in order to prevent illegitimate address space use.

4.3 Address Markets

Given that free address pools are now mostly exhausted and that demand for IPv4 address space will likely continue to grow (at least until significantly broader IPv6 deployment), address space transfers arise as a natural step necessary to further distribute address space to those networks that need it. In light of the issues discussed above—namely the fragmentation of addresses into legacy and non-legacy address blocks subject to varying RIR-policies, and connected ownership discussions, as well as the lack of widely adopted resource certification mechanisms—the landscape of such address transfers is at best murky. Network operators have already started buying and selling address blocks under varying conditions, as we outline in the following. This resulted in the emergence of several *address brokers* (e.g., [1, 48, 49]); companies that assist network operators wanting to buy or sell address space. Eventually, the RIRs learned to encourage the use of address brokers to mediate transactions within the strict confines of RIR policies.

4.3.1 RIR Transfer Policies

Today, four[15] out of the five RIRs allow address space transfers among their members [12, 14, 54, 83]. In addition, ARIN, RIPE and APNIC offer *transfer listing services* [8, 24, 75], where network operators can list address blocks they want to sell and express the need for certain amounts of address space they want to buy. These services aim to help interested parties to come together, but use of them is not mandatory. RIPE publicizes aggregated statistics for address space requests and offerings, listing fewer than one million available addresses, and more than 17 million requested addresses, as of February 2015 [75]. These listings do not include any prices, as the negotiations remain entirely at the discretion of the respective parties.

We observe a striking difference between how the RIRs perceive their roles when it comes to conducting transfers under their policies. Except for RIPE, the RIRs still require the receiving party of a transfer to *justify* their need for more address space according to their already established policies. For example, APNIC requires transfer recipients to document use rates for past allocations, as well as detailed plans for the use of transferred resources [12], while ARIN states that recipients must demonstrate the need for up to a 24-month supply following their established policies [14].

RIPE—as of February 2014—removed all "justification of need" clauses from their policies. Address space can be transferred from any member to any other member without the need to make statements of how the transferred addresses will be used by the recipient. The proposal [82] argued that address conservation will be in the interest of the members themselves (to not waste address space). With regard to concerns about possible address hoarding by wealthy LIRs it states that "*markets* [for other commodity goods] *function well and in a competitive manner, and there is no reason why the trade of IPv4 addresses will be any different*".

As of June 2014, *Inter-RIR Transfers*—i.e., transfers between address holders in different regions—are only possible between ARIN and APNIC. ARIN explicitly requires justification of need on the receiving side of a transfer—even if the recipient is located in a different region [17]. Thus, RIPE's removal of justification of need in its transfer policies rules out Inter-RIR transfers with the ARIN region.

Figure 6 summarizes the transfer policies in place by the RIRs along with the address space they administer.

Another scenario in which address block transfers happen—and happened long before modern transfer policies were established—is due to *Mergers & Acquisitions*. In this case, address blocks are part of the assets of a company. Since the related contracts are often confidential, these transfers are not publicly listed by the RIRs—with the exception of APNIC, which requires full disclosure of the involved parties and publicly lists the corresponding address blocks [7]. The RIR's documents make no explicit statements about the justification of need for the transferred allocations. ARIN only states that the transferred resources will be subject to ARIN policies [14], while APNIC states that it will "review the status" of the allocations, requiring full disclosure of all allocations held by the "entities in question". If that is not provided, APNIC will "require that they be returned" [12].

4.3.2 Transfers Outside the RIRs

Given that neither the legal nor the technical aspects of address space transfers are under the full control of the current RIR framework, parties can also conduct transfers separately from the RIRs. To the extent that these occur, a definitive determination of the party possessing a given allocation becomes more difficult because the RIRs no longer possess accurate records.

Even though ARIN states that legacy holders are subject to ARIN policies, recent transfers, such as the well-known sale of more than 660K IP addresses from the Nortel bankruptcy to Microsoft, have raised concerns whether

[13] ARIN requires legacy resource holders to sign an LRSA in order to be eligible to register their resources in the RPKI. Moreover, ARIN requires any operators wanting to *use* the ARIN RPKI data to sign a Terms of Service Agreement that includes an indemnification clause [23].

[14] To overcome this, AS-Path validation is necessary [57].

[15] AFRINIC states the possibility of transfers between LIRs [3], but prohibits any such transfers in their LSA unless they arise due to Mergers & Acquisitions [2].

Juniper Exhibit 1006, page 7

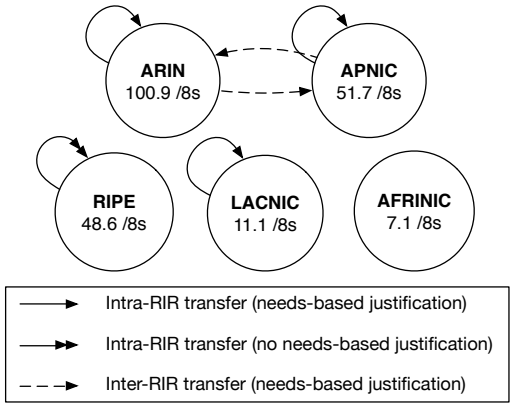

Figure 6: Current address space transfer policies of the RIRs and the administered address space.

they complied with proper ARIN transfer policy. Mueller et al. [59] state that while ARIN was formally involved in the transfer, likely no needs-based evaluation was performed on the receiving side, and that ARIN's intervention boiled down to a *"face-saving exercise"*. As the relationship of legacy holders towards the RIRs is not entirely clear, one IP address trader has suggested that legacy address holders in the ARIN region could de-register their address space there and re-register it with a different RIR, such as RIPE [50]. Doing so would effectively allow inter-region transfers from ARIN to RIPE without undergoing any transfer process. But currently there is no process to de-register from an RIR.

Aside from transfers that were formally noticed by the RIRs (such as the above example), address transfers can also happen without the involvement of any registry at all. While address space can be of various types (*legacy*, allocated to a holder by an RIR, assigned by a holder to an end-user, PI-assigned directly from the RIR to an end-user), bound to various contractual limitations, not much prevents any party from unofficially transferring an address block to another entity. This is known as a *"black market"* transfer. This possibility stems from the decoupled nature of address block management and actual address block use. If RIRs do not acknowledge such transfers, registry information becomes in turn inaccurate and incomplete, making the attribution of address blocks to their respective holders difficult.

In the simplest terms, we can view a transfer as simply an address block—or parts of it—formerly in use by some entity A now being used by some entity B, possibly outside the purview of any RIR regulation. If the routing of the concerned address block is possible after the transfer (it is not filtered by networks), and (to a lesser degree) the corresponding reverse-DNS zones become under the control of the receiver (e.g., by subdelegation of reverse-DNS zones by the previous owner), the transfer would be successful.

It is unclear whether it is even feasible to detect the occurrence of such transfers. Livadariu et al. attempted to detect such transfers by looking for changes in routing origins over time [58]. One difficulty here is that transferred address blocks are not necessarily routed before they are transferred. Indeed, prior routing might be unlikely, as unrouted address space is likely also unused and thus more likely to be transferred. Also, whether such a transfer would be reflected in the reverse-DNS is unclear, as NS records might simply not be changed and PTR records might be unchanged or switched off. Shifts in traffic, latency changes or geographical changes might be due to transfers but also due to restructurings within a company.

Thus, defining the boundaries of what exactly an address transfer is and what it is not is not straightforward. It is likely that the official RIR transfer policies only cover a fraction of the total address transfers occurring in various instantiations of the above scenarios. While transfers undergoing the RIRs policies are publicly listed [7, 16, 76] and quantifiable, the number of address transfers outside this framework is unknown and requires further research.

5. OVERCOMING SCARCITY

IP address scarcity has become reality. That is, today only ARIN and AFRINIC still hand out address space under regular conditions, while RIPE's, APNIC's and LACNIC's pools have become exhausted, and they only hand out one small allocation from their last /8 to a requesting LIR. Comparing allocation rates in 2014 to allocation rates in previous years, it is clear that the supply of address blocks from the RIRs cannot satisfy the demand. Thus, address shortage problems require other approaches. Generally, we can consider three possible solution spaces for this problem: (i) develop more address space by adopting IPv6, (ii) multiplex current IPv4 address space using address sharing techniques such as Carrier-grade NAT (CGN), and/or (iii) more efficiently use the current IPv4 address space.

Develop more address space. The successor to IPv4, IPv6 [35], extends the routable address space by orders of magnitude. (Its design also aimed to address some other shortcomings of IPv4, such as support for mobility and extensibility.) It reflects the ultimate natural solution to the scarcity problem. The RIRs advocate its use (e.g., RIPE hands out remaining IPv4 address blocks only to LIRs that have already received an IPv6 allocation [83]), and the community has undertaken many other efforts to promote IPv6 adoption (e.g., [47]). Nevertheless, the fraction of both IPv6-enabled networks as well as native IPv6 traffic on the Internet remains comparably small—adoption of IPv6 remains problematic and only slowly increases.[16] IPv6 is by itself not compatible with IPv4, and requires complex transition mechanisms to ensure compatibility between the IPv4 and IPv6 Internet (e.g., [63]).

Multiplex address space. Alternatively, we can get by with many fewer addresses by multiplexing. Enterprise networks have long employed NAT to avoid having to allocate individual public IP addresses to every Internet-attached device. Today, numerous approaches to perform address sharing at scale are available—see [85] for a comprehensive study—and are already in use by several large ISPs. While widespread use of NAT raises concerns about eroding end-to-end connectivity and semantics, as well as concerns by law enforcement agencies due to the erosion of attribution of IP addresses to end-users [37], it poses fewer compatibility issues than IPv6 when employed for legacy network infrastructure. According to [26], already more than 3% of

[16] As of February 2015, Google reports some 4.5% of clients accessing Google to be IPv6 enabled, with adoption rates as high as 28% in Belgium, around 10 to 15% in the US and Germany, and increasing support in other European countries. Nonetheless, the per-host adoption rate still ranges at or below 1% for most countries, including China, India and Russia [42].

Juniper Exhibit 1006, page 8

Internet users are behind CGNs, and Web hosting companies already employ heavy address sharing.

Use address space more efficiently. As visible in Figure 4, about a third of all Internet address blocks remain unrouted, and thus not in (at least public) use. Moreover, even routed address space is not necessarily in active use. As mentioned above, recent studies find utilization levels for the routed address space at around 50% to 60% [33, 34, 89]. Hence, significant usable address space remains. Making more efficient use of address space will require adapting address management policies, guidelines and technologies, including the difficult (both technically and politically) problem of re-assigning already allocated address blocks.

Network operators are currently adopting all of these options to varying degrees: IPv6 adoption, CGN, and address transfers. We would expect that cost will determine the manner and timeline when different options predominate. We should in turn find these costs reflected in the price of IPv4 addresses as exchanged via secondary markets.

6. OUTLOOK

IPv4 address scarcity is an issue requiring attention from the networking and research community. Depending on the success of transitioning towards widespread use of IPv6, we face a mid-to-long term scenario in which IPv4 addresses will have significantly more demand than supply. The key question is for how long the cost of IPv4 addresses will be viewed as lower than the cost of transitioning to IPv6 or using CGN.

That said, we note that while the limited IPv4 address space clearly will not suffice in the longer term to provide every Internet device with its own address, the current scarcity arises due to address management practices, and not (yet) due to protocol limitations. Large fractions of the address space remain unrouted, and of those address blocks that are routed, again only a fraction is actually in use.

Just how to adapt the governance of the available address space to the current situation remains a pressing question. While it is unclear whether IP address block holders have ownership rights for their IP addresses, secondary markets already exist to facilitate their exchange. However, the uncertainties associated with address space transfers—both the legal status of legacy address blocks and the varying policies among RIRs when it comes to such transfers—will also complicate how pricing develops. This, in turn, makes it increasingly difficult for network operators to make decisions on which technology to adapt when.

As IPv4 addresses become an ever more scarce resource, increasing numbers of transfers, both inside the RIR framework as well as outside, are likely. As transfers outside the RIR framework can result in less accurate registration data provided by the RIRs—which in turn limits the possibility to use formal defenses such as RPKI-based origin validation—address block hijacking events presumably will also increase. Viewing IP addresses as resources, other issues arise, such as resource certification and the exercise of control over "who uses what address space". As an inherently global resource, it is questionable whether the distributed registry framework can cope with the looming issues and provide sufficient resource liquidity. Future scenarios for the management could include a more competitive environment among RIRs, or even a re-centralization of the registries.

From a research perspective, several issues arise: How to overcome scarcity issues? What technologies will have what impact on the Internet? Will the community succeed in fully deploying IPv6 within the next decade, or will we find ourselves stuck in a long-term situation in which IPv4, IPv6 and technologies like CGN operate in parallel? What will be the corresponding impact on the Internet topology, its performance, and its reliability?

How to effectively deploy resource certification of address blocks and how to ensure routing only by the respective holder? How commonly do address transfers occur outside the RIR framework and with what sort of historical development and likely future trends? What measurements could inform recommendations on how to govern the address space, in light of both IPv4 and IPv6 allocations? Did the creation of the distributed registry framework influence topological properties? How should the RIRs agree on implementing consistent policies?

We argue that the Internet community as a whole would greatly benefit from empirical studies tackling the above questions, which will both aid network operators with resolving business-critical decisions, as well as policy makers as they adapt to this new landscape and work towards ensuring further unhindered growth of the Internet.

Acknowledgments

This work was partially funded by NSF grant CNS-1111672 and DHS/ARL contract W911NF-05-C-0013.

7. REFERENCES

[1] Addrex. IPv4 Address Broker. `http://www.addrex.net`.

[2] AFRINIC. AFRINIC Service Agreement 2013. `https://www.afrinic.net/en/services/rs/rsa`.

[3] AFRINIC. IPv4 Allocation Policy (AFPUB-2005-v4-001). `http://www.afrinic.net/en/library/policies/126-afpub-2005-v4-001`.

[4] AFRINIC. Out-Of-Region Use of AFRINIC Internet Number Resources (AFPUB-2014-GEN-002-DRAFT-01). `http://afrinic.net/en/community/policy-development/policy-proposals/1157-out-of-region-use-of-afrinic-internet-number-resources`.

[5] AFRINIC. Policy Development Process in the AFRINIC service region (AFPUB-2010-GEN-005). `http://www.afrinic.net/en/community/policy-development/251-policy-development-process-in-the-afrinic-service-region-afpub-2010-gen-005`.

[6] APNIC. APNIC IPv4 Address Pool Reaches Final /8 (2011-04-15). `http://www.apnic.net/publications/news/2011/final-8`.

[7] APNIC. IPv4 Address Transfer Logs. `ftp://ftp.apnic.net/public/transfers/apnic/`.

[8] APNIC. IPv4 Transfers Listing Service. `http://www.apnic.net/services/become-a-member/manage-your-membership/pre-approval/listing`.

[9] APNIC. Policy development process (APNIC-111). `http://www.apnic.net/publications/media-library/documents/policy-development/development-process`.

[10] APNIC. Policy environment for Internet number resource distribution in the Asia Pacific (APNIC-125,

Juniper Exhibit 1006, page 9

2011-05-09). https://www.apnic.net/policy/policy-environment/text.
[11] APNIC. prop-017: Recovery of unused address space (2004-02-26). https://www.apnic.net/policy/proposals/prop-017.
[12] APNIC. Transfer, merger, acquisition, and takeover policy (APNIC-123). http://www.apnic.net/policy/transfer-policy.
[13] ARIN. ARIN Policy Development Process. https://www.arin.net/policy/pdp.html.
[14] ARIN. ARIN's Number Resource Policy Manual (NRPM) Version 2014.02. https://www.arin.net/policy/nrpm.html.
[15] ARIN. Draft Policy ARIN-2014-1: Out of Region Use. https://www.arin.net/policy/proposals/2014_1.html.
[16] ARIN. Inter-RIR and Specified Transfers of Internet Number Resources. https://www.arin.net/knowledge/statistics/transfers.html.
[17] ARIN. Inter-RIR Transfers. https://www.arin.net/resources/request/transfers_8_4.html.
[18] ARIN. Legacy Registration Service Agreement Statistics. https://www.arin.net/knowledge/statistics/legacy.html.
[19] ARIN. Legacy Registration Services Agreement. https://www.arin.net/fees/agreements/legacy.html.
[20] ARIN. Legacy Registration Services Agreement Outreach (2009-03-03). https://www.arin.net/resources/legacy/outreach.html.
[21] ARIN. Legacy Registration Services Agreement v3.0. https://www.arin.net/resources/agreements/legacy_rsa.pdf.
[22] ARIN. Registration Services Agreement v11.0. https://www.arin.net/resources/agreements/rsa.pdf.
[23] ARIN. RPKI Terms of Service Agreement. https://www.arin.net/resources/rpki/tos.pdf.
[24] ARIN. Specified Transfer Listing Service. https://www.arin.net/resources/transfer_listing/.
[25] ARIN. Transfers to Specified Recipients. https://www.arin.net/resources/request/transfers_8_3.html.
[26] I. Van Beijnum. With the Americas running out of IPv4, its official: The Internet is full (2014-06-12). http://arstechnica.com/information-technology/2014/06/with-the-americas-running-out-of-ipv4-its-official-the-internet-is-full/.
[27] I. Van Beijnum. *BGP: Building reliable networks with the border gateway protocol*. Page 61ff. O'Reilly Media, Inc., 2002.
[28] R. Bush. Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI). RFC 7115 (Best Current Practice), January 2014.
[29] R. Bush, R. Austein, S. Bellovin, and M. Elkins. The RPKI & Origin Validation. NANOG 52, 2001.
[30] K. Butler, T. Farley, P. McDaniel, and J. Rexford. A Survey of BGP Security Issues and Solutions. *Proceedings of the IEEE*, 98(1), 2010.
[31] V.G. Cerf. IAB Recommended Policy on Distributing Internet Identifier Assignment and IAB Recommended Policy Change to Internet "Connected" Status. RFC 1174 (Informational), 1990.
[32] J. Czyz, M. Allman, J. Zhang, S. Iekel-Johnson, E. Osterweil, and M. Bailey. From Infancy to Adolescence: A Decade-long Measurement of IPv6 Adoption. In *ACM SIGCOMM*, 2014.
[33] A. Dainotti, K. Benson, A. King, K. Claffy, E. Glatz, and X. Dimitropoulos. Estimating Internet Address Space Usage Through Passive Measurements. *ACM CCR*, 44(1), 2014.
[34] A. Dainotti et al. Lost in Space: Improving Inference of IPv4 Address Space Utilization. Tech. rep., CAIDA, Oct 2014. http://www.caida.org/publications/papers/2014/lost_in_space/.
[35] S. Deering and R. Hinden. Internet Protocol, Version 6 (IPv6) Specification. RFC 1883 (Proposed Standard), 1995. Obsoleted by RFC 2460.
[36] W. Deng, P. Zhu, X. Lu, and B. Plattner. On Evaluating BGP Routing Stress Attack. *Journal of communications*, 5(1):13–22, 2010.
[37] Donald Codling, FBI. Attribution Growing Challenges For LEAs. CWAG Conference of Western Attorneys General 2013 http://www.cwagweb.org/pdfs/2013/privacy_conference_ppts/cwag%20preso%2013.ppt.
[38] P. Ford, Y. Rekhter, and H. Braun. Improving the Routing and Addressing of IP. *IEEE Network*, 7(3), 1993.
[39] V. Fuller, T. Li, J. Yu, and K. Varadhan. Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy. RFC 1519 (Proposed Standard), 1993. Obsoleted by RFC 4632.
[40] Geoff Houston. IP Address Report. http://www.potaroo.net.
[41] E. Gerich. Guidelines for Management of IP Address Space. RFC 1466 (Informational), 1993. Obsoleted by RFC 2050.
[42] Google. IPv6 Statistics. https://www.google.com/intl/en/ipv6/statistics.html.
[43] K. Hubbard, M. Kosters, D. Conrad, D. Karrenberg, and J. Postel. Internet Registry IP Allocation Guidelines. RFC 2050 (Historic), 1996. Obsoleted by RFC 7020.
[44] ICANN. Global Policy for the Allocation of the Remaining IPv4 Address Space (ratified in 2009). https://www.icann.org/resources/pages/remaining-ipv4-2012-02-25-en.
[45] ICANN. Multi-Stakeholder Discussion: Legacy Internet Protocol (IP) Numbers in the Current Policy Environment (2012). Audio Transcript. http://toronto45.icann.org/node/34377.
[46] ICANN Blog. Recovering IPv4 Address Space. http://blog.icann.org/2008/02/recovering-ipv4-address-space/.
[47] Internet Society. World IPv6 Launch. http://www.worldipv6launch.org.
[48] IP Trading. IPv4 Address Broker. http://www.iptrading.com.
[49] IPv4 Market Group. IPv4 Address Broker. http://www.ipv4marketgroup.com.
[50] IPv4 Market Group Blog. Don't Sign an RSA During Your 8.2 IPv4 Transfer. http://ipv4marketgroup.com/?s=do+not+sign+rsa.

Juniper Exhibit 1006, page 10

[51] IRR. List Of Routing Registries. http://www.irr.net/docs/list.html.

[52] A. Khan, K. Hyon-chul, T. Kwon, and Y. Choi. A comparative Study on IP Prefixes and their Origin ASes in BGP and the IRR. *ACM CCR*, 43(3), 2013.

[53] LACNIC. Policy Development Process (version 2.0 - 01/09/2008). http://www.lacnic.net/en/web/lacnic/proceso-de-desarrollo-de-politicas.

[54] LACNIC. Policy Manual v2.1 (2014-03-25). http://www.lacnic.net/web/lacnic/manual-2.

[55] LACNIC. Registration Services Agreement. http://lacnic.net/docs/rsa-en.pdf.

[56] M. Lepinski and S. Kent. An Infrastructure to Support Secure Internet Routing. RFC 6480 (Informational), February 2012.

[57] M. Lepinski and S. Turner. An Overview of BGPSEC. IETF Internet-Draft draft-ietf-sidr-bgpsec-overview-05 https://tools.ietf.org/html/draft-ietf-sidr-bgpsec-overview-05, 2014.

[58] I. Livadariu, A. Elmokashfi, A. Dhamdhere, and K. Claffy. A First Look at IPv4 Transfer Markets. In *ACM CoNEXT*, 2013.

[59] M. Mueller, B. Kuerbis, and H. Asghari. Dimensioning the Elephant: An Empirical Analysis of the IPv4 Number Market. *info*, 15(6):6–18, 2013.

[60] National Science Foundation (NSF). Amendment 7 to Cooperative Agreement Between NSI and U.S. Government (1997-12-03). http://archive.icann.org/en/nsi/coopagmt-amend7-03dec97.htm.

[61] RIPE NCC. YouTube Hijacking: A RIPE NCC RIS case study. http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study.

[62] Network World, Vol. 17, No. 4. Stanford move rekindles 'Net address debate (2000-01-24).

[63] E. Nordmark and R. Gilligan. Basic Transition Mechanisms for IPv6 Hosts and Routers. RFC 4213 (Proposed Standard), October 2005.

[64] NRO. Extended Allocation and Assignment Reports. http://www.nro.net/statistics.

[65] NRO. Free Pool of IPv4 Address Space Depleted. http://www.nro.net/news/ipv4-free-pool-depleted.

[66] NTT America. Request for IPv4 and IPv6 Address Space v3.7. https://whois.gin.ntt.net/ipj.txt.

[67] Packet Clearing House (PCH). Routing Origin Inconsistency. https://prefix.pch.net/applications/routing-origin-inconsistency/.

[68] J. Postel. Assigned numbers. RFC 790 (Historic), 1981. Obsoleted by RFC 820.

[69] J. Postel. Internet Protocol. RFC 791 (Internet Standard), 1981. Updated by RFCs 1349, 2474, 6864.

[70] A. Ramachandran and N. Feamster. Understanding the Network-level Behavior of Spammers. *ACM CCR*, 36(4), 2006.

[71] Y. Rekhter, B. Moskowitz, D. Karrenberg, and G. de Groot. Address Allocation for Private Internets. RFC 1597 (Informational), March 1994. Obsoleted by RFC 1918.

[72] J.K. Reynolds and J. Postel. Assigned numbers. RFC 1060 (Historic), 1990.

[73] RIPE NCC. Early Registration Transfer (ERX) Project. http://www.ripe.net/lir-services/resource-management/erx.

[74] RIPE NCC. Ingrid Wijte on the RIPE policy mailing list. (2012-05-23). http://www.ripe.net/ripe/mail/archives/address-policy-wg/2012-May/006981.html.

[75] RIPE NCC. IPv4 Transfer Listing Service. http://www.ripe.net/lir-services/resource-management/listing.

[76] RIPE NCC. IPv4 Transfer Statistics. http://www.ripe.net/lir-services/resource-management/ipv4-transfers/table-of-transfers.

[77] RIPE NCC. Policy Development Process in RIPE (ripe-614). http://www.ripe.net/ripe/docs/ripe-614.

[78] RIPE NCC. RIPE NCC Standard Service Agreement (ripe-533). http://www.ripe.net/ripe/docs/ripe-533.

[79] RIPE NCC. RPKI Statistics. http://certification-stats.ripe.net.

[80] RIPE NCC. Services to Legacy Internet Resource Holders (ripe-605). http://www.ripe.net/ripe/docs/ripe-605.

[81] RIPE NCC. Status of Legacy IPv4 Address Space (2011-09-12). https://labs.ripe.net/Members/xavier/status-of-legacy-ipv4-address-space.

[82] RIPE NCC. Post Depletion Adjustment of Procedures to Match Policy Objectives, and Clean-up of Obsolete Policy Text. http://www.ripe.net/ripe/policies/proposals/2013-03, 2013.

[83] RIPE NCC. IPv4 Address Allocation and Assignment Policies for the RIPE NCC Service Region (ripe-606). http://www.ripe.net/ripe/docs/ripe-606, 2014.

[84] E. M. Rubi. The IPv4 Number Crisis: The Question of Property Rights in Legacy and Non-Legacy IPv4 Numbers. *American Intellectual Property Law Association (AIPLA) Quarterly Journal*, 39:477, 2011.

[85] N. Skoberne, O. Maennel, I. Phillips, R. Bush, J. Zorz, and M. Ciglaric. IPv4 Address Sharing Mechanism Classification and Tradeoff Analysis. *IEEE/ACM ToN*, 22(2), 2014.

[86] F. Solensky. Continued Internet Growth. In *IETF 18*, pages 59–61, 1990. http://www.ietf.org/proceedings/18.pdf.

[87] Andree Toonk. How accurate are the Internet Route Registries (IRR). (2009-03-28). http://www.bgpmon.net/how-accurate-are-the-internet-route-registries-irr/.

[88] Andree Toonk. Securing BGP routing with RPKI and ROA's. (2011-01-19). http://www.bgpmon.net/securing-bgp-routing-with-rpki-and-roas/.

[89] S. Zander, L. Andrew, and G. Armitage. Capturing Ghosts: Predicting the Used IPv4 Space by Inferring Unobserved Addresses. In *ACM IMC*, 2014.