The Great IPv4 Land Grab: Resource Certification for the IPv4 Grey Market

Eric Osterweil
eosterweil@verisign.com

Shane Amante
Shane.Amante@Level3.com

Danny McPherson
dmcpherson@verisign.com

Dan Massey
massey@cs.colostate.edu

ABSTRACT

The era of free IPv4 address allocations has ended and the grey market in IPv4 addresses is now emerging. This paper argues that one cannot and should not try to regulate who sells addresses and at what price, but one does need to provide some proof of ownership in the form of resource certification. In this paper we identify key requirements of resource certification, gained from both theoretical analysis and operational history. We further argue these requirements can be achieved by making use of the existing reverse DNS hierarchy, enhanced with DNS Security. Our analysis compares reverse DNS entries and BGP routing tables and shows this is both feasible and achievable today; an essential requirement as the grey market is also emerging today and solutions are needed now, not years in the future.

1. INTRODUCTION

The era of freely available IPv4 addresses is over, but the demand for IPv4 addresses has not ceased. Prior to February 3rd, 2011, any organization could request IPv4 addresses from one of five regional registries, and after meeting some technical requirements, the organization could obtain IPv4 address space at a nominal cost. As these requests depleted the pool of unallocated space at regional registries, they turned to the Internet Assigned Numbers Authority (IANA) for additional unallocated space. On January 31, 2011 the APNIC regional registry made one such request, which depleted the pool of available IPv4 space at IANA. This triggered the February 3rd distribution of the five remaining /8 unallocated blocks to the regional registries[4, 5]. As a result, the regional registries can no longer receive additional IPv4 address blocks from the IANA. In other words, we have run out of unallocated IPv4 addresses.[1]

What was essentially a free resource is now a scare commodity and this has profound implications for the Internet, as many organizations appear to still favor IPv4 over IPv6 allocations. The exhaustion has changed a fundamental property of the Internet, and simple economics dictate that an IPv4 market must now exist. Unable to obtain IPv4 address space from the regional registries, organizations must find other means and will turn to what are currently grey markets in address space. Even the popular press has reported headlines such as "Microsoft Spends 7.5 Million Dollars on IP Addresses" [6]. It is not a question of whether a market for IPv4 addresses will exist, it is a question of what is needed to keep the IPv4 land grab from devolving into a maelstrom of address swindling. The question for the technical community is what characteristics can and/or should we try to impose to safeguard transactions in such a market?

Consider the potential problems where an organization believes they have purchased IPv4 addresses for their web site, begins using them, and then gets sued by an ISP who claims they *own* a covering prefix. This paper argues the fundamental technical question is proof of ownership and we draw an analogy to the real estate market. Regulators do not determine who can buy and sell real estate and they do not set or otherwise approve the purchase price for property. They do record the transactions and provide title to the property. Further, the title itself is not guaranteed and, at least in the U.S., buyers are required by lenders to purchase title insurance to protect against the rare cases where the title is imperfect. Without some notion of a property title (e.g. proof of ownership), one could simply point to a portion of land and attempt to sell it to an unsuspecting victim. Similarly, without some proof ownership, one can simply point to an IP address block and sell it to an unsuspecting victim. The old classic con of "I can offer you a great deal on the Brooklyn Bridge" may now be replaced with "I can offer you a great deal on 10/8". Clearly, one should not purchase the Brooklyn Bridge or 10/8, but would you be interested in a deal for 129.82.138/24?[2] Any market place in addresses without proof of ownership is fundamentally flawed.

The regional registries cannot dictate terms of sales,

[1]Technically, IANA has run out of available IPv4 space; but the registries are also quickly running through their space.

[2]This address is assigned to one of the authors by the legitimate owner, but the author is explicitly not authorized to sell, barter, or otherwise exchange this address space.

but would simply like to be informed when transfers happen. The process of allocation and assignment is non-trivial. Both the process and the current market are discussed further in Section 2. While we believe registries should not be in the position of approving sales, how could one hope to provide proof of ownership given complex assignments where not all transactions will be reported to registries? We contend that the fundamental challenge is providing a near-term solution that certifies the ownership and policies of IPv4 resources, and aligns operational costs and benefits.

We answer this challenge in Section 3 by first setting the requirements for a resource certification system that can essentially prove ownership, similar to a title in a real estate transaction. We then show that such a system is not only achievable, but we use the reverse DNS as the basis for a system that can be deployed immediately, securely, and incrementally by interested parties in today's Internet. Whenever an allocation or assignment is made from a registry, the owner is already granted access to the corresponding reverse address space in DNS, today. In fact, this is common practice as allocations are sub-allocated from registry to registry to ISP. This is done for operational reasons such as the practical need for mail relays to have reverse DNS entries, but our interest in reverse DNS is that it is already an existing practice that IP owners maintain because it aligns administrative costs with operational benefits. Thus even in a grey market there is an operational motivation to keep the reverse DNS in some minimal working order. Further, a feasibility study in Section 5 shows that 94.3% of reverse zones are in a position to deploy the approach today and *all zones* could deploy by following existing standard practices set forth in Internet RFCs.

2. THE IPV4 ADDRESS MARKET

The market for buying and selling allocations of IPv4 space has begun to shape up in ways that many people expected [21, 19, 11]. It exists outside the boundaries, policies, and control of registries, or any other form of explicit governance.

2.1 Registry Policies Today

Blocks of IPv4 addresses (*netblocks*) are recursively allocated from a pool held by a global root *registry* down to end users. Initially, the root entity, called the Internet Assigned Numbers Authority (IANA), is responsible for the entire pool of IPv4 addresses (as well as other resources, but here we focus IPv4 addresses). IANA's role is to *allocate* blocks of these addresses to the set of five Regional Internet Registries (RIRs): ARIN, RIPE, APNIC, LACNIC, and AfriNIC. After an allocation, the receiving RIR becomes responsible for that portion of the global pool. Similarly, other entities (such as ISPs, local registries, etc.) will apply to one of the RIRs for *allocations* or *assignments* of IPv4 address blocks. The RIR will then sub-allocate portions of their *free pool* to the client, and that client can then use or subdivide these blocks further. At the final step in this recursive process, the resource holder will *assign* a portion of their available IPv4 block to a company for usage. The mechanism by which the ARIN registry tracks changes that happen lower in the recursive process is the Shared Whois Project (SWIPs), but even this runs on the honor-system.

A policy distinction exists between blocks of IP addresses that have been *allocated* and blocks that have been *assigned*. Specifically, assigned blocks are not allowed to be sub-assigned. [10, 17, 20, 2, 8] Thus, the assignment policy for 129.82.138/24 is expected to restrict its sale implicitly. Just as when someone leases access to land, they are allowed to occupy and control it, but are restricted from selling it.

2.2 The Market Today

Perhaps the most publicized instance of the IPv4 grey market was Microsoft's acquisition of 666,624 IPv4 addresses from Nortel. After entering into bankruptcy, Nortel began selling off its assets, which included its pool of increasingly coveted IPv4 addresses. Specifically, Nortel sold 666,624 of them for $7.5 million to Microsoft. Based on this trade alone, the going rate for IPv4 addresses would be roughly $11.25.[6]

In addition to large public examples like this one, some people have sought to set up online "market places" to broker acquisitions. For example, several organizations [3, 1, 7] have already stepped forward to act as IP-brokers.

2.3 Certification, not Regulation of Re-Allocation

Faced with the fact that the grey market is already a reality, registries are in no position to dictate the rules of how IPv4 addresses will be sold. For example, should ARIN feel entitled to legally block Microsoft's business transactions with Nortel? Further, as trades take place farther and farther outside of the public stage, is it even possible for the registries to *know* when a sale happens? Clearly, the answer to both of these is "no." In fact, this setting casts registries as *subordinates* to entities participating in the grey market! This is because the registries lose the ability to control their address pools as soon as they allocate them, and thus are not authorities over their addresses any more. In fact, ARIN has publicly stated that they will not try to interfere with the grey market. However, they maintain that if they find themselves in a position to verify a sale that has violated their policies, they believe they can take action. However, participants need a way to provide *resource certification* without implicitly requiring regulation of

the grey market.

3. REQUIREMENTS AND PRINCIPLES

In order for a resource certification framework to address the needs and problems of the IPv4 grey market, its design must meet correctness requirements and must *also* be realistic enough to be deployed and maintained in the Internet's chaotic setting.

3.1 Correctness Requirements

Must be publicly checkable and open: In order to benefit an open market place, the certification of IPv4 resources must allow arbitrary parties to verify ownership of those resources. This is analogous to the way in which one would engage a title company to verify the ownership of land. Indeed, the distributed administration model of the Internet suggests that since anyone can engage in a transaction with anyone else, the function of a title company should not be an exclusive right of certain parties, but must be open.

Must disallow transfers to multiple parties: If a seller owns an address block, he/she must not be able to sell it to multiple buyers. The resource certification system must provide timely information to show transactions have occurred and must not allow multiple simultaneous transfers.

Reflects, but does not predetermine the result of market transactions : The resource certification framework should not try to authenticate the rights or nature of sales, or it surely would face the same obstacles that the registries now face. Rather, this framework must simply reflect the *results* of transactions.

Must allow registries to reflect their policies: In addition to the grey market transactions that assign and re-assign ownership, a resource certification framework must also capture important policy distinctions that the registries make about IPv4 blocks. Specifically, if a registry has assigned a block, that block is not allowed to be sub-allocated. The registry's policy is an important part of this equation, but potentially distinct from the actions in the market place. That is, many expect that policy prescriptions will prove to be too restrictive, and will become obsolete.

If the system fails to meet any of the above requirements, it simply does not solve the desired resource certification problem. At the same time, meeting the above requirements does not mean it will provide an effective deployable solution. We next discuss several additional requirements that we argue are essential if the system is to actually succeed in the Internet.

3.2 Deployment Requirements

Owners of IP real estate must be able to manage their own certifications: Owners of IPv4 addresses must have the same level of autonomy in managing their allocations and assignments as any other operation they perform on the Internet.

Must be incentive based, not mandate based: Mandates have a poor record in the Internet. Best practices are routinely ignored, obsolete systems and versions continue to operate, attempts to remove problematic aspects and attributes from protocols are simply ignored. If one could easily enforce mandates, perhaps one could simply mandate IPv4 must be replaced by IPv6. We take as a given that one cannot mandate use, but one can provide incentives, and/or align operational costs with benefits.

Returning to our real estate example, one cannot prevent two parties from striking a private deal to buy land and never updating the title to the property. The solution to this problem is not mandates, it is incentives. If the new buyer ever decides to sell the property, having a title makes it easier to sell to wider range of potential buyers and thus increases the likely value of the property. In addition, how useful would buying land be if the new owner could not turn on utilities because they could not prove they were the rightful owner? The seller also has incentives to transfer title; such as formally moving taxes and other liabilities to the new buyer. Similarly, we argue that address blocks with valid resource certification are more valuable since they can be more easily traded and resource certification can help place liability with the proper owner. The tangible benefits of having proper certification should far out weigh the costs of creating and maintaining the certification.

It must be deployable today! The IPv4 market already exists, so any relevant solution must be feasible in today's setting, and it must be realistic to see deployment in the very near-term. Unfortunately, this means that there is no time for a green-field approach because such an undertaking would likely be overtaken by the evolution of the market place and could become obsolete while still being designed.

4. REVERSE DNS SOLUTION

Our premise is that one can use the *existing DNS* and the reverse DNS in particular. We make *no changes* to the DNS servers that provide data and *no changes* to the DNS resolvers that request data. Further, unlike some previous approaches to adding BGP routing information to the reverse DNS, we also make *no changes* to the reverse DNS structure. This choice is dictated by both our requirement to deploy today and our requirement to be incentive driven. Modifying servers or resolvers would take time and requires a large incentive for operators to invest in the new system. We instead leverage the existing reverse DNS tree that operators must already maintain. Registries already follow RFC 2050 [14], which mandates that when they allocate IP space, they must also delegate control of a correspond-

ing branch of the reverse DNS. That is, when IANA allocated the block 129/8 to ARIN, it also delegated `129.in-addr.arpa`. When ARIN allocated 129.82/16 to a university, it also delegated `82.129.in-addr.arpa`, etc. Thus, every owner of an allocation has the ability to add records to the corresponding portion of DNS. Operationally, this has become critical for many reasons. First, due to the incidence of email spam, many large ISPs refuse to accept email connections from servers or clients that don't have DNS `PTR` records in the reverse DNS that correspond to an incoming mail server's (or client's) IP address. This is done because only the rightful owner of that IP space can put a record in the right zone for that IP. Further, new record types, such as `SSHFP`, are becoming increasingly popular for the same reason: administrative control of the reverse zone. In fact, a new IETF working group [16] is investigating this general approach further.

The *one addition* we make to this existing structure are two new record types. First, we propose to describe the Address Policy in a new type called `ADDRPOL`. This record will simply reflect whether a netblock is an allocation or an assignment. Its position in an authoritative zone describes what the owner claims, and its position in the parent (or covering) zone describes the policy when it was delegated. The other type is intended to disambiguate how the netblock is being used and by whom. The new `SRO` type is used to describe the list of Autonomous System Numbers (ASNs) that an owner would like to authorize as valid origins for IPv4 addresses within the netblock. Finally, the addition of DNSSEC [9] allows one to authenticate the result.

4.1 Handling Classful Allocations

The reverse DNS is delegated on octet boundaries and ideally suited for classful allocations (e.g. mask lengths of 8, 16, 24, or even 23). We map a classful allocation to a DNS name in a straight-forward manner; the allocation mask length is denoted by `m`⟨*masklength*⟩ and the prefix is written in the standard reverse DNS notation *using the minimum number of octets needed to define the allocation*. For example, 129/8 is mapped to the DNS name `m8.129.in-addr.arpa` and 129.82/16 is mapped to the DNS name `m16.82.129.in-addr.arpa`. Any site interested in the ownership and policy associated with 129.82/16 can simply query for the RRset (`m16.82.129.in-addr.arpa,IN,ADDRPOL`) or (`m16.82.129.in-addr.arpa,IN,SRO`). Note as discussed above, no change is needed to resolvers, servers, or the reverse DNS structure.[3]

The restriction to using the minimum number of octets is essential to provide a unique name and also essential in directing the query to appropriate reverse DNS zone. For example, 129.82/16 could also be written as 129.82.0.0/16 and even as 129.82.1.2/16. Note the mask length of 16 implies only the first two octets (16 bits) are necessary to define the prefix. The last 16 bits are unnecessary and they direct queries to the wrong zone.

To understand this, suppose organization A has been allocated 129.82/16 and, following RFC 2050 [14], also operates the `82.129.in-addr.arpa` zone. Further suppose organization A has assigned 129.82.0/24 to organization B. Again following RFC 2050, the reverse DNS zone `0.82.129.in-addr.apra` should be delegated to organization B. The name `m16.82.129.in-addr.arpa` belongs to the `82.129.in-addr.arpa` zone and is run by organization A. Organization A can set the appropriate `ADDRPOL` and `SRO` entries and sign these entries using the `82.129.in-addr.arpa` DNSSEC key. However, the name `m16.0.0.82.129.in-addr.arpa` has too many labels, the `0.0` labels cause this name to belong to the `0.82.129.in-addr.arpa` zone, which is run by organization B. Organization B would need to set the appropriate `ADDRPOL` and `SRO` entries and sign these entries using the `0.82.129.in-addr.arpa` DNSSEC key owned by organization B. By using the minimum octets, one avoids giving the assignee (organization B) authority and responsibility to set the allocation and ownership policy of organization A's address block.

4.2 Handling Classless Allocations

Often overlooked in DNS proposals is handling the classless allocations. In other words, one should not overlook an allocation such as 129.82.0/18. Such allocations and assignments are clearly commonplace.The approach to solving this can make the difference between a novel idea and deployable system. Our design adds an additional constraint in that we assume one cannot add any new structures. Resolvers, servers, and the reverse DNS tree are fixed elements in our design. Note, however, that reverse DNS must handle such allocations anyway. Otherwise, the owner of 129.82.0/18 would not be able to operate email servers. Thus, we leverage the *existing system* to store our `ADDRPOL` and `SRO` records.

To do so, we first need to extend the naming scheme to classless allocations. We again denote the mask length as `m`⟨*masklength*⟩, require using the minimum number of octets needed to define the allocation, and use the standard reverse DNS notation. However, the mask length must be inserted as *second label* (from left to right). Thus 129.82.0/18 is written as `0.m18.82.129.in-addr.arpa`. The placement of the mask length as the second label is essential to placing the authority for the `ADDRPOL` and `SRO` records at the correct zone. Figure 1 depicts this conversion.

[3]Technically one could encode the desired data in an `TXT` record and not even need to introduce the `ADDRPOL` or `SRO` record types, but modern resolvers and servers support unknown RR types and as authors of DNS related RFCs, we encourage using valid types over overuse/abuse of `TXT`.

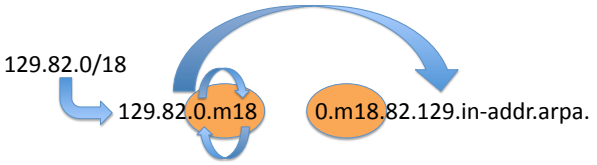

Figure 1: How to convert a /18 into a DNS zone.

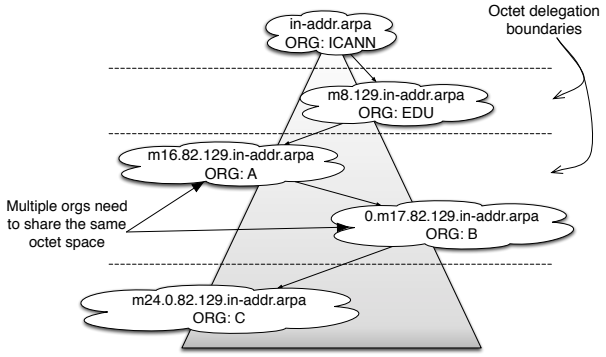

Figure 2: The m⟨*masklength*⟩ technique and RFC 2317 allow us to split authority between octets.

Note there is nothing unique about the mask length of 18. The same reasoning applies to any mask length between 17 and 23. This works for any mask length.

4.3 Handling Classless Allocations With Multiple Organizations

A final challenge occurs when classless allocations and assignments are done at multiple layers. To illustrate the problem, we return to our example where organization A has been allocated 129.82/16. Further organization A has allocated 129.82.0/17 to organization B, and B has further assigned 129.82.0/24 to organization C. We first note that reverse DNS should be implemented as well. The reverse zone `82.129.in-addr.apra` should be delegated to organization A and `0.82.129.in-addr.apra` should be delegated to organization C. In addition, organization B should be the authority for all entries corresponding to 129.82.0/17, but organization A does not want to share administration of its zone `82.129.in-addr.apra` zone (and especially its DNSSEC zone key) with organization B. Thus, our approach uses the m⟨*masklength*⟩ technique to share authority between octets, as seen in Figure 2.

Specifically, RFC 2317 [12] offers the standardized best practice for how this should be accomplished. Organization A creates CNAME entries (e.g. DNS aliases) for all the entries belonging to organization B. In this case, organization A will create a CNAME for `0.82.129.in-addr.apra`, `1.82.129.in-addr.apra`, ..., `127.82.129.in-addr.apra`.[4] There may be better ways to handle such delegations; however RFC 2317 is the *standardized* way to achieve this. Previous work in this area [18] also proposed another new technique for classless delegation in reverse DNS, and like RFC 2317, our approach would be able to ride on top of that mechanism if it were to be certified and standardized by the IETF. In short, our approach is compatible with the standard and other known approaches that support recursive delegations, but full discussion is omitted due to space restrictions.

5. HOW MUCH WORK IS NEEDED?

One very large question from the design in Section 4 is, "how much effort is needed to deploy this framework today?" The answer to this question comes in two parts. First, we need to understand how much of the reverse DNS is properly delegated to the rightful owners of IPv4 netblocks. These represent the easy low hanging fruit; these netblock owners simply need to insert the resource certification records in their zones and deploy DNSSEC. Second, how much of the reverse DNS is *not* properly delegated and would require additional steps to comply with the RFCs and/or update to resolve undesired administrative assignments. In this case, either reverse zones *could* be delegated that would allow netblock owners to manage a corresponding DNS zone, or netblocks are allocated (or assigned) multiple times between octet boundaries, and the conventional usage of reverse DNS does not suffice. We define this last case to be an *ownership conflict*.While this does not limit the ability of netblock owners to use this approach, it does constitute an additional step to deploy. In order to gauge the incidence of this complication, we examined the structure of today's reverse DNS and use BGP routing announcements to represent current IPv4 allocations and assignments.

We began by measuring how much of the reverse IP space (at or above the /24 mask) maps to separate DNS zones on May 8^{th}, 2011. Specifically, we queried for `SOA` records for the zones ranging from `0.in-addr.arpa` to `255.255.255.in-addr.arpa` (corresponding to 0/0 to 255.255.255/24). We chose to use BGP to provide an approximation of IPv4 ownershipand examine a BGP routing table from a tier-1 service provider. Based on previous findings [13], we believe that this RIB is representative of the active prefixes throughout the Internet.

We found that this portion of the reverse DNS had approximately 3.78 million unique DNS zones, and the BGP routing table was comprised of 349,295 BGP prefixes announced from 37,226 unique ASNs. Some large organizations announce their BGP prefixes from multi-

[4]Names with additional labels such as `0.0.82.129.in-addr.apra` do not need to be added. Delegating `0.82.129.in-addr.apra` is sufficient See RFC 2317 for details.

ple (different) ASNs. Therefore, using a mapping table [15], we correlated ASNs to organization names. Thus, by mapping prefixes to the reverse zones and their organizations we were able to count the number of organizations that *can* manage their reverse DNS without a conflict, and how many organizations are in conflict.

Of the 3.78 million zones, BGP only needed 137,180 reverse zones (3.63%) to map all of its prefixes. Initially, 129,296 reverse zones (94.3%) did not contain conflicts, based on prefix alone (i.e. not even considering organization mapping). These represent netblock/zones that could be administered by a single organization without needing RFC 2317 (i.e. `1.m24.2.3.in-addr.arpa`).

However, today many branches under the reverse DNS do not extend far below the first label (`x.in-addr.arpa`, equivalent to x/8). To enhance the expressiveness of the reverse DNS tree in order to meet the RIB's needs, 201,839 extra (more specific) zones were added. We note that this operational addition would be needed for any use of the reverse DNS, and therefore represents a general operational benefit, not solely for the resource certification framework. This increased the number of non-conflicting zones to 191,320 (96.2%). Moreover, when using organizational mapping information, we found that the number of zones that are conflict-free is really 195,645 (98.4%). This means that 98.4% of the Internet can deploy this resource certification framework today, with a trivial amount of effort! This leaves conflicts in 3,280 zones, who are still fully able to use RFC 2317 to obtain their proper delegation.

6. CONCLUSION

Certain aspects of meat-space land grabs illustrate important lessons for the Internet. In 1889, Harper's Weekly published commentary on the Oklahoma land grab. Noting that no resource management considerations were made beforehand, Harpers characterized the chaos of the event as being the result of "massive stupidity of federal policy." IPv4 has been on the verge of "running out" for years, but it has now happened to the IANA free-pool. While RIRs may still have addresses, they are rationing their remaining allocations and they will eventually run out as well. This makes proper management of all IPv4 addresses a critical issue. Rather than face the sorts of swindling and chaos that could ensue from an uncertified pool of sold, resold, subdivided, re-resold, and further sub-divided IPv4 netblocks, the Internet needs some form of order.

In this work, we have shown that by using the reverse DNS, one can add a very critical element to crystallize *resource certification* in a way that aligns operational costs with benefits. Our approach is not only incrementally deployable, but *does not* require any infrastructure changes on either the authorities source or the clients.

Our analysis has shown that even though there are allocation scenarios that complicate the deployment of this framework: 1) 98.4% of the current Internet zones do not need to alter their deployment of the reverse DNS zones, and 2) those that *do* have complications can resolve them in RFC compliant ways, and they need to anyway in order to run services such as email servers.

Finally, we believe this approach opens up many new avenues for secure systems in the Internet, such as: IPv6 resource certification, IPv6 neighbor discovery, incrementally deployable routing security [22], and more.

7. REFERENCES

[1] Addrex. `http://www.addrex.net`.
[2] APNIC guidelines for IPv4 allocation and assignment requests. `http://www.apnic.net/`.
[3] Depository. `http://www.depository.net`.
[4] IANA IPv4 Address Space Registry. `http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xml`.
[5] IPv4 IANA Free Pool Depletion FAQ. `https://www.arin.net/resources/request/ipv4_depletion.html`.
[6] Microsoft spends $7.5m on ip addresses. `http://www.theregister.co.uk/2011/03/24/microsoft_ip_spend/`.
[7] tradeipv4: Ipv4 address block exchange. `http://www.tradeipv4.com`.
[8] AfriNIC. Policy - IPv4 Address Allocation Policies. `AFPUB-2005-v4-001`.
[9] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. DNS Security Introduction and Requirement. RFC 4033, March 2005.
[10] ARIN. ARIN Number Resource Policy Manual. `https://www.arin.net/policy/nrpm.html`.
[11] Network World Carolyn Duffy Marsan. Need IPv4 addresses? Get 'em here. April 25, 2011.
[12] H. Eidnes, G. de Groot, and P. Vixie. Classless IN-ADDR.ARPA delegation. RFC 2317, March 1998.
[13] Kaustubh Gadkari, Daniel Massey, and Christos Papadopoulos. Dynamics of prefix usage at an edge router. In *PAM'11*, 2011.
[14] K. Hubbard, M. Kosters, D. Conrad, D. Karrenberg, and J. Postel. Internet Registry IP Allocation Guidelines. RFC 2050, Nov 1996.
[15] Geoff Huston. List of ASN organizations. `http://www.potaroo.net/bgp/iana/asn.txt`.
[16] IETF. DANE. `https://datatracker.ietf.org/wg/dane/charter/`.
[17] LACNIC. LACNIC - Policies. `http://lacnic.net/en/politicas/index.html`.
[18] Ang Li, Xin Liu, and Xiaowei Yang. Bootstrapping accountability in the internet we have. NSDI'11, 2011.
[19] Milton L Mueller. Scarcity in IP addresses: IPv4 Address Transfer Markets and the Regional Internet Address Registries. *Internet Governance Project Paper IGP08002*, (September 1981):1–20, 2008.
[20] RIPE. IPv4 Address Allocation and Assignment Policies for the RIPE NCC Service Region. `RIPE-509`.
[21] Ars Technica Iljitsch van Beijnum. ARIN fights IP address trading as transition to IPv6 may get new deadlines. August 2008.
[22] Xiaowei Yang and Xin Liu. Internet Protocol Made Accountable. In *HotNets-VIII*, 2009.