Analysis of the IPv4 Address Space Delegation Structure

Anusha Sriraman Kevin R.B. Butler Patrick D. McDaniel Padma Raghavan
Department of Computer Science and Engineering
Pennsylvania State University
{sriraman, butler, raghavan, mcdaniel}@cse.psu.edu

Abstract

The Internet has grown tremendously in terms of the number of users who rely on it and the number of organizations that are connected to it. Characterizing how this growth affects its structure and topology is vitally important to determine the fundamental characteristics and limitations that must be handled, such as address space exhaustion; understanding the process of allocating and delegating address space can help to answer these questions. In this paper, we analyze BGP routing data to study the structure and growth of IPv4 address space allocation, fragmentation and usage. We explore the notion of delegation relationships among prefixes and use this information to construct an autonomous system (AS) delegation tree. We show that delegation in the Internet is not significantly correlated to the underlying topology or AS customer-provider relationships. We also analyze the fragmentation and usage of address space over a period of five years and examine prefixes that are delegated by organizations vs. those that are not delegated. We notice that the address space usage due to delegating prefixes is increasing at the same rate as the address space usage due to non-delegating prefixes. This indicates that fragmentation rate of the address space is actually almost a constant with respect to total address usage. Additionally, we show that most delegation is performed by a small number of organizations, which may aid in the implementation of a public-key infrastructure for the Internet.

1. Introduction

The Internet is a system in perpetual flux. The number of Internet users has grown prodigiously over the past decade, and the number of organizations leveraging the network has correspondingly increased. Understanding how the tremendous expansion of the Internet affects its structure and topology is critical to predicting its evolution and determining fundamental characteristics of its growth, as well as potential fundamental limitations that must be grappled with.

Individual organizations collectively comprise the global Internet. These organizations form administrative domains called *autonomous systems* (ASes), and receive address space signifying ownership of that space on the Internet. This space is advertised by the peers of an AS and is circulated throughout the Internet, letting other entities know that this space belongs to a particular organization. The Border Gateway Protocol (BGP), provides the means for ASes to advertise their address space, and in turn, for other ASes across the world to reach these addresses. Address space is allocated to individual organizations through a process of *delegation* that can be traced back to international numbering agencies and ICANN [1], the global authority for addressing and naming for the Internet. By understanding the ways that blocks of addresses, known as *prefixes*, are allocated, it is possible to determine an underlying *delegation structure* of the Internet.

The aim of this paper is to study the structure and growth of the Internet based on the delegation structure of the IPv4 address space. ASes can delegate a part or all of their prefixes to other ASes. We represent the address space as a tree of routable prefixes constructed on the basis of such delegation relationships. This prefix tree provides an insight into relationships, such as business relationships, between organizations.

We examined the BGP routing table data of the over 40 route servers that contribute to the Route Views routing data repository, and inferred delegation patterns based on this information. With this data, we derived the AS delegation tree and compared it with existing results that infer topology and relationships between customers and providers, to determine what correlations exist. It is commonly thought that when an organization becomes the customer of a provider organization (e.g., an Internet Service Provider, or ISP), the provider organization will allocate a portion of their address space to the customer for their use. This is an easier method for the customer to receive address space, as they do not have to perform the administrative work of obtaining that space from international and global address registries, and do not have to configure their networks to make this

1 ICANN is the International Consortium of Assigned Names and Numbers.

address space routable. However, our analysis determined that a large fraction (almost 90%) of ASes that exhibit a provider-customer relationship do not share an address delegation relationship.

Additionally, we performed a longitudinal study spanning five years of routing tables, and processed over 1.6 billion table entries to study the growth and fragmentation of the IPv4 address space, and how the space delegated by ASes evolved compared to the non-delegated address space. We defined a notion of *delegating prefixes*, which are fragmented into smaller prefixes that are further delegated and advertised, and *non-delegating prefixes*, which are not further delegated. We observed that the address space covered by delegating prefixes is increasing at around the same rate as the address space covered by non delegating prefixes. Furthermore, the rate of new prefixes seen in the Internet is approximately equal between delegating and non-delegating prefixes; thus, the fragmentation rate of the available address space is almost constant.

The results of these analyses are notable in their demonstration of the constant behavior of delegation and acquisition of new address space on the Internet. For example, although economic and business conditions have changed dramatically since 2001, the rate at which prefixes are allocated to organizations has been virtually constant, indicating that the collective Internet has operated in a manner independent of external conditions and may be expected to continue operating in a similar manner. Additionally, these results validate previous work into the nature of delegation as described by McDaniel et al. [20]. Even over a long-term period, there is relatively little correlation between topology and delegation; most delegation is done by a small number of organizations. From a security standpoint, this is an important result for mitigating the complications of deploying a public-key infrastructure (PKI) in the Internet.

The remainder of this paper is organized as follows. Section 2 presents the related work. Section 3 gives a background on BGP and AS graphs. We also describe the data formats, extraction and filtering. Section 4 explains delegation structures and methodologies used to construct them. In Section 5 we analyze and correlate the AS graphs. This section also introduces the study and analysis of address space fragmentation and the related results. Section 6 presents discussion of the results and Section 7 concludes the paper.

2. Related Work

A significant amount of work has been done to understand and model the Internet's topology and to determine an underlying structure. While some studies focused on determining connectivity through AS topology graphs [8, 9, 16, 26], other research has found that a power-law topology exists within the Internet [11, 22, 24]. These results are based on the examination of information such as routing updates and analysis of traffic. Our approach considers the routing tables but determines structure over a longer period of time, measured in years. Additionally, while these studies concentrate on topology, our work considers the address delegation relationships and the corresponding structure of the address space that is created as a result.

Information from the BGP routing tables has been extensively used to determine characteristics of both Internet routing and the actual structure of the Internet. Gao devised an algorithm for inferring contractual AS relationships from BGP routing tables, classifying relationships between customers, providers, peers, and siblings and using heuristics to form these classifications [15]. Subramanian et al. [19] expanded on this idea by using multiple vantage points throughout the Internet and using metrics to determine the accuracy of these observations. Tools such as Rocketfuel [28] use routing data to accurately perform tomography, or determining the internal characteristics of ASes.

Many researchers have been actively studying the BGP routing data to understand evolution and growth of the address space. Broido et al. studied the growth of prefixes and ASes in [5, 6, 7]. Huston provides extensive reports on BGP routing data in [12, 13, 14]. He models the address distribution of IANA and RIR's and extrapolates this information to predict the exhaustion of IPv4 address space. Address allocation patterns and its impact on routing table growth were studied in [17, 23, 29]. Our analysis differs from these as we explore the notion of addresss delegations (introduced in [20]) to study the allocation and evolution of the address space. We used address delegations to construct and analyze the delegation hierarchy in the Internet. We spanned this study over a period of 5 years with a large number of data points. We processed around 1.6 billion advertisements overall. To the best of our knowledge, this is the first detailed study of delegation hierarchy in the Internet.

3. Background

3.1. BGP Data: Acquisition and Extraction

The Internet consists of large numbers of interconnected hosts called autonomous systems (ASes). The Border Gateway Protocol [30] is the interdomain routing protocol used on the Internet. Each BGP speaking router in an AS advertises *address prefixes* of varying length, which represent a block of address space owned by the AS. Prefixes are advertised using CIDR notation in a manner such as $a.b.c.d/x$ where $x \leq 32$ and the number of addresses represented by this prefix is $2^{32-x} - 2$, such that a prefix of form $a.b.c.d/24$ represents 254 addresses.[2] These prefixes are advertised by an AS to its peers, who in turn send this information to their neighbors, thus allowing reachability information to spread throughout the entire network. The *AS path* represents the

2 Typically, the first and last addresses of a block are the identifier and broadcast addresses and are not part of the usable space.

route through a set of ASes required for a given AS to send packets to a desired destination prefix. Since BGP is a *path-vector* protocol, the metric used to determine the optimal path to a destination is generally the shortest path, or number of ASes that require traversal. However, ASes very often use policy considerations to optimize and bias the routing process, based on contractual and peering agreements with customers, providers, and competitors. The resulting routes selected may thus be very different. Each BGP-speaking router stores the best route to a destination in its routing table. If a new optimal route is found, the router advertises that route to its peers and updates its routing table accordingly.

Over the years, researchers have extensively used BGP routing tables to study the routing patterns, growth and topology of the Internet. Many tools and services are available to collect BGP routing data. In our analysis we used the repository of BGP routing table information collected by the University of Oregon's Route Views project [10]. Route Views archives BGP tables from a set of over 40 listening points from around the world. BGP data from Route Views is available in different types of data formats. For our study, we used BGP RIB (Routing Information Base) files, which represent snapshots of each peer's routing table collected every 2 hours.

We collected the BGP tables at an interval of every 5 days from October 2001 to June 2006, retrieving one routing table per day (1400 EST) . For each table, we extracted a unique set of routable prefixes $\{P\}$ and the associated set of ASes $\{A\}$ along with the AS–prefix mappings, $\{P\} \leftrightarrow \{A\}$. We also extracted all AS paths advertised in the table, $\{APath\}$. While extracting data from the tables, we removed all advertisements related to private ASes (AS numbers in the range 64512-65535), prefixes of mask lengths greater than /24, which are typically the smallest blocks that appear in the routing table, as major ISPs often filter smaller blocks to prevent an explosion of entries in the routing table. [3]

RouteViews collects BGP data from many peers. We used only those prefix advertisements common to most of the peers. Additionally, ICANN keeps a fixed list of /8 addresses in the space between 0.0.0.0 and 127.255.255.255 (the former "Class A" space) that includes reserved space; we filtered advertisements that were misadvertised as originating in these reserved spaces.

3.2. AS Graphs

Because of the decentralized nature of the Internet, it can be difficult to determine the relationships individual ASes possess with each other and what the collective sum of these relationships looks like. An AS graph is a method of showing the topology of the Internet by representing ASes as nodes and their connections to other Ases as connectivity links. There can be different kinds of AS graphs based on the type of connectivity information used to construct the graph. ASes share different kinds of relationships based on geographical proximity, logical connectivity, and business agreements. The following is a description of AS graphs used in this paper.

- **AS Connectivity Graph**
  An AS connectivity graph [8, 9, 16] is constructed from the set of AS paths, $\{APath\}$. This graph represents the physical AS connectivity in the Internet. We constructed this graph by extracting the ASes and their associated edges from the routing data. If the AS path is $a\ b\ c$, we place a connectivity link between ASes a & b and ASes b & c. Over the years, a large number of new ASes have been added to the Internet which has caused the number of such AS links to grow exponentially.
- **AS Provider–Customer Graph**
  Gao [15] designed an algorithm to infer contractual provider–customer AS relationships from routing tables. This algorithm parses the AS path and identifies the AS with highest degree as the top provider. Consecutive AS pairs (a, b) to the left of the top provider are assigned a transit relationship in which AS b provides transit services to AS a. Consecutive AS pairs (a, b) to the right of the top provider are assigned a transit relationship in which AS a provides transit services to AS b. Provider-customer relationships can then be inferred based on the assumption that a provider will provide transit services for its customer but the customer will not provide any transit service to the provider. For example, in (a, b), a is a provider to the customer b if a provides transit services to b and b does not provide transit services to a; thus, (a, b) forms a provider-customer relationship.The detailed algorithm, provided in [15], has been widely accepted as a method for extracting AS relationships and we use it in our paper to construct the provider-customer graph of ASes.
- **AS Delegation Tree**
  ASes own certain ranges of address blocks called prefixes. An AS can delegate all or a portion of its prefix to another AS. On doing so, these ASes establish a business relationship. We define this as a *delegation relationship*. If AS a delegates a portion of its prefix to AS b, ASes a and b have a delegation relationship with a being the delegating AS and b being the customer AS. This customer AS can keep this prefix or further delegate it to other customers and so on. This presents a hierarchical structure of prefix delegations among ASes. Using this notion of delegation relationships between ASes, we construct the AS delegation tree. We present

3 Traditionally, tier-1 ISPs would often filter advertisements smaller than a /19 prefix, but the growth of customers with smaller routable address spaces has led to /24 blocks now being the smallest routable prefixes.

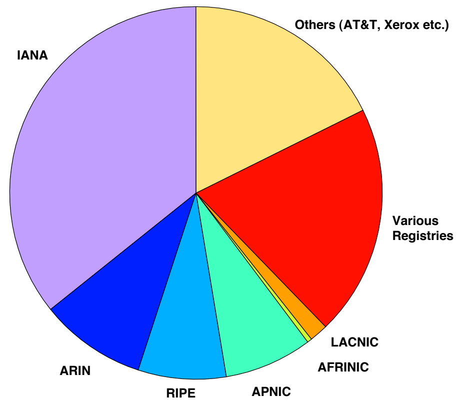

Figure 1: Address Allocation to Registries and ISPs.

details about the construction of this tree in the following section.

4. Constructing the Delegation Tree

For a given set of prefixes $\{P\}$, we infer the relationship between prefixes in the following manner. A prefix a is said to be delegating prefix b if the address space covered by prefix b is a subset of the space covered by a, e.g. prefix $12.0.0.0/8$ is delegating prefix $12.1.0.0/16$. Prefix a is then said to be the *delegating prefix* of prefix b. A prefix can be related to many other prefixes, but it has only one direct parent delegating prefix, e.g., $12.1.0.0/16$ is a parent of $12.1.1.0/24$, but $12.0.0.0/8$ is also a parent of $12.1.1.0/24$. We resolve this ambiguity with the *longest prefix matching* rule: we infer the delegating parent prefix to be the one with the highest subnet mask. Hence in our example, we infer that $12.1.0.0/16$ is the delegating parent prefix of $12.1.1.0/24$. This type of relationship between prefixes is called *prefix delegation*.

Using the notion of prefix delegations described above, we constructed a hierarchical delegation structure of the address space, or *prefix delegation tree*. The prefix delegation tree was constructed as follows. The root of this tree was chosen to be $0.0.0.0/0$ which covers the entire address space. The second level in this tree has the nodes $0.0.0.0/1$ and $128.0.0.0/1$. These prefixes are not advertised as they do not exist in the routing tables and were added only to provide a structure to the delegation tree. IANA[4] manages the entire IPv4 address space [3]. Hence, IANA is the root of the prefix delegation tree (associated with the prefix $0.0.0.0/0$). IANA delegates some portions of the address space directly to organizations and some to Regional Internet Registries (RIRs). These allocations are in chunks of /8 prefixes. The largest unit of address advertisement is a /8 prefix and hence, these /8 prefixes form the next level of the

4 The operation of IANA, the Internet Assigned Numbers Authority, is managed by ICANN, and we refer to each synonymously.

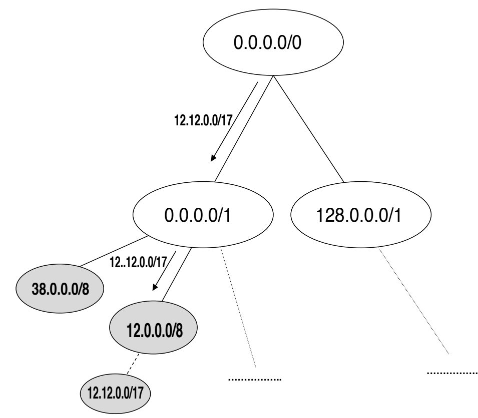

Figure 2: Construction of the prefix delegation tree.

delegation tree. The RIRs then delegate the address space assigned to them to other organizations and Local Internet Registries (LIRs).

IANA allocates address space to various RIRs such as ARIN, RIPE, APNIC, AfriNIC, LACNIC etc. We obtained the /8 allocations from IANA [4], as some of these /8's are not advertised pubilcly via BGP. Figure 1 shows the address allocation to IANA and other registries. 46.6% of /8's are allocated to RIRs and 17.7% of /8's are allocated to other organizations that obtained their space directly from IANA (AT&T, Xerox etc.). 35.7% of /8's is reserved by IANA for either future delegations or for other purposes, e.g., multicast and private addresses.

To construct the remaining tree, we used the prefixes obtained from BGP routing tables. We introduce a prefix into the delegation tree by finding the parent delegating prefix and adding a delegating link from this parent to the child prefix. We find this parent prefix by traversing the prefix delegation tree from the root along the path of related prefixes until we find the immediate delegating parent of this prefix. Figure 2 shows the addition of prefix $12.12.0.0/17$ to the prefix delegation tree.

We defined the *AS delegation tree*, which depicts the delegation of prefixes between ASes, in Section 3. To construct the AS delegation tree, we used the prefix delegation tree and mapping $\{P\} \leftrightarrow \{A\}$ to find the corresponding AS for every prefix tree node. Most prefixes are origin stable [25] and are associated with a single AS, hence we can assume that this mapping is fairly stable. In AS delegation tree, an AS may be associated with multiple tree nodes depending on the number of prefixes advertised by it.

5. Analysis Results

We analyzed BGP data from Oct 2001 to June 2006, collecting the BGP tables at an interval of 5 days. For each table, we extracted and filtered the data as described in Section 3, and constructed the AS connectivity and AS provider-customer graphs. Using the prefix set $\{P\}$ and the

Table 1: Characteristics of the prefix delegation tree.

Prefix Type	Fraction
Delegating Prefixes $\{DP\}$	8.3%
Non Delegating Prefixes $\{NDP\}$	91.7%
Top Prefixes $\{TP\}$	49.5%
Top Delegating Prefixes $\{TP\} \cap \{DP\}$	5.7%
Top Non Delegating Prefixes $\{TP\} \cap \{NDP\}$	43.8%
Fragments	50.5%

methodology described in Section 4 we constructed prefix and AS delegation trees. To further the understanding of prefix delegations and address space fragmentation, we define the following.

- Top prefixes: The set of prefixes $\{TP\}$ that are delegated directly from /8 prefixes. $\{TP\}$ represents prefixes that are *allocated* by RIRs or top ISPs that in turn obtain their address space directly from IANA. These prefixes appear at the fourth level in the delegation tree. The first and second levels are occupied by IANA and the third level is occupied by /8 prefixes (RIRs and ISPs).
- Delegating prefixes: The set of prefixes $\{DP\}$ that are further fragmented and delegated to other ASes. These prefixes appear as *internal* nodes in the tree.
- Non-delegating prefixes: The set of standalone prefixes $\{NDP\}$ that are are not delegated any further. These prefixes appear as *leaf* nodes in the tree.
- Fragments: All prefixes other than the top prefixes and /8's are called fragments and represented as $fragments = \{P\} - \{TP\}$. Fragments are the result of fragmentation and delegation of existing prefixes.
- Address usage: Address usage by a prefix is the fraction of the total address space covered by this prefix. e.g., address usage by $12.0.0.0/8$ is $\frac{2^{24}}{2^{32}}$.

We studied the evolution of the prefix delegation tree with time to note address delegation and allocation patterns. Table 1 lists some properties of the tree. Analysis reveals that almost 92% of the nodes in the tree are leaf nodes, i.e., non-delegating prefixes. A much smaller fraction, 8%, of the prefix nodes are delegating prefixes. This shows that a very small fraction of prefixes are involved in delegation, and supports the notion that prefix delegations are highly centralized.

We observed that 49.5% of the prefix nodes are top prefixes, i.e., almost half of the advertised prefixes are allocated by RIRs and top ISPs. These top prefixes are further fragmented and delegated, thereby generating the delegation structure. Out of these top prefixes, 5.5% of the prefixes are delegating prefixes and the remaining 44% are non-delegating prefixes. A large fraction of non-delegating top prefixes signifies that a large number of end users or customers obtain their address space directly from the top delegators and RIRs. 90% of the prefixes are seen either as the

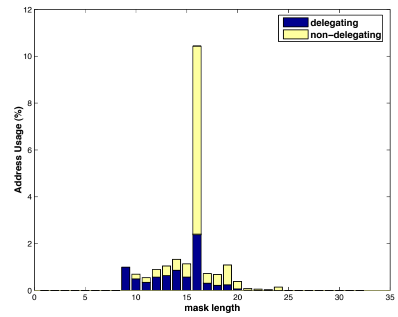

(a) Address space usage, 25 December 2001.

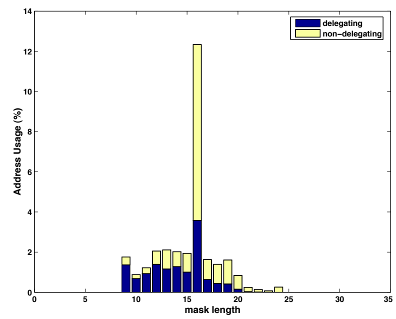

(b) Address space usage, 26 April 2006.

Figure 3: Address space usage of Top Prefixes, considering *delegating* and *non-delegating* prefixes.

top prefixes or as prefixes delegated directly from top prefixes. Although the height of the tree is 7, only 10% of the delegations extend to this length. We observed that these quantities were not subjected to much change in the study period. This indicates a fairly stable growth pattern of the address space.

Figure 3 shows the address usage of different mask lengths in the top prefixes. We notice that the address usage is dominated by /16s. Also, the usage of delegating prefixes is greater than that of non-delegating prefixes for mask lengths lesser than 16. For mask lengths greater than 16, the delegating prefix usage is lesser than the non delegating prefix usage. Further, we noticed that this pattern is preserved through the entire span of study, as evidenced by the similarities between Figures 3(a) and 3(b), which differ by over 4 years. This supports the notion that larger prefixes are more likely to be fragmented than smaller prefixes.

5.1. Analysis of AS Graphs

To understand how prefix delegations relate to the Internet topology and existing business models, we compared the AS connectivity and AS provider-customer graphs with the AS delegation tree. We make this comparison by using two types of links – subtree links and sibling links. Subtree links correspond to ancestor-descendant links in the tree, and signify delegation relationships. Sibling links are links between ASes that obtain prefixes from the same parent del-

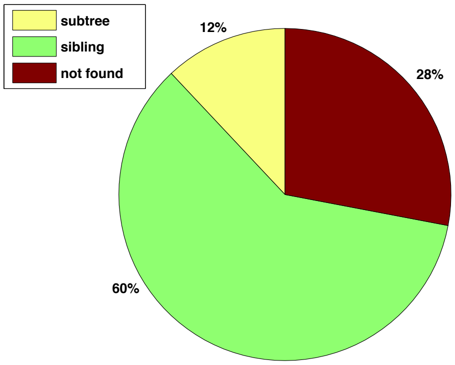

Figure 4: Comparison of the AS provider-customer graph with the delegation tree.

egating prefix. For every link a–b in the AS connectivity graph and provider-customer graph, we recorded if the corresponding link was present in the delegation tree. If a occurred in the subtree of b, we recorded a subtree link between a and b. If a and b share a common parent in the delegation tree, i.e., if they are siblings in the delegation tree, we recorded a sibling link between a and b.

We noticed that around 60% of the links in the AS provider-customer graph were found to be siblings in the delegation tree. A small fraction, 12% of these links were found to share delegation relationship as subtree links in the tree(Figure 4). For both AS graphs (topological connectivity and provider-customer), around 28% of the links were not present in the tree at all. These statistics were found to be consistent with time from 2001 to 2006. Almost 90% of these sibling links are related at the top level (among top prefixes) in the tree. It can be inferred from this that most ASes that are related through sibling links get their address space from RIRs or large ISPs. We can hence conclude that there is no significant correlation between the AS topology or the AS provider-customer graphs and the AS delegation tree. We verified our results by using the provider-customer graphs inferred by CAIDA [2] and found them to be consistent.

5.2. Address Space Fragmentation and Usage

Address space usage may be determined by identifying the top prefixes. The address space consumed by these prefixes indicate the actual usage of the Internet space. RIRs and ISPs are allocated with /8 address blocks, and in turn allocate space to other ISPs and end users. These allocated prefixes appearing in $\{TP\}$ are either fragmented and delegated further, or kept for internal use. In either case, we consider this address block as *used*. We view the growth in address usage by top delegating prefixes ($\{DP\} \cap \{TP\}$) and top non-delegating prefixes ($\{NDP\} \cap \{TP\}$). Figure 5 shows the address usage by /8 allocated space, top prefixes, and delegating and non-delegating top prefixes. We observed that the address space usage due to delegat-

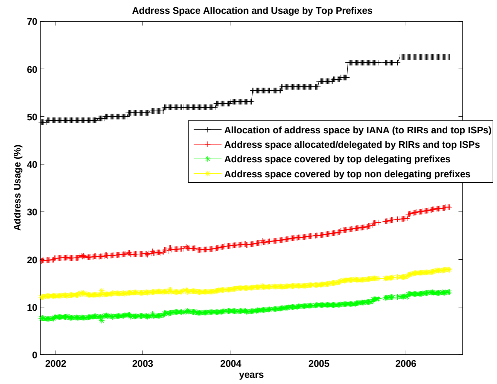

Figure 5: Address Space Usage by top prefixes: Allocation by IANA, top delegating and top non-delegating prefixes.

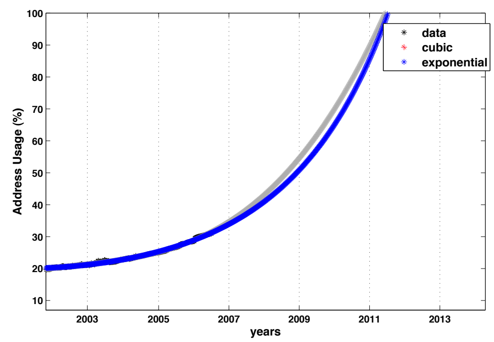

Figure 6: Modeling: Address Space Usage by top prefixes.

ing top prefixes and non-delegating top prefixes are growing approximately at the same rate. This shows that both address allocation and fragmentation are increasing at the same rate. We also note that the address usage by the non-delegating top prefixes is higher than the address usage by delegating top prefixes. This indicates a large fraction of direct allocation to customers who do not participate in further delegations.

We modeled the address space usage of top prefixes. Using curve-fitting methods, we found that the data supported either a cubic or exponential growth rate; these curves are shown in Figure 6. We extrapolated both curves that fit the model and observed that the exhaustion point for the top prefixes will occur in approximately 2012. At this exhaustion point, the total address space covered by $\{TP\}$ is 100% and no further allocations from the top ISPs and RIRs will be possible; any further growth of the address space will be due to fragmentation of the existing space.

Figure 7 shows the growth in the number of advertised prefixes. We observe that the number of non-delegating prefixes are growing at a much faster rate than the number of delegating prefixes. Fragments represent prefixes that result from delegation and fragmentation of the allocated address space. Figure 8 shows the growth in the counts of top prefixes and fragments. We observe that the allocated prefixes (top prefixes) and fragments have similar rates of growth.

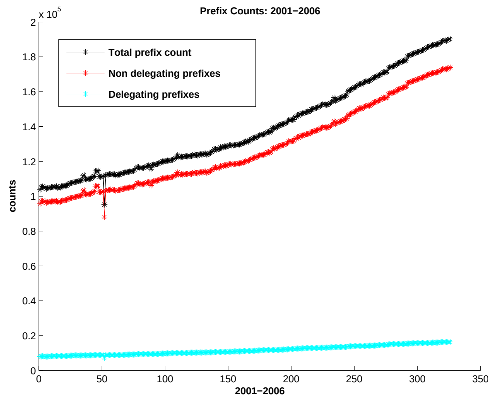

Figure 7: Prefix Counts: Total, Delegating, Non–delegating and Top prefixes.

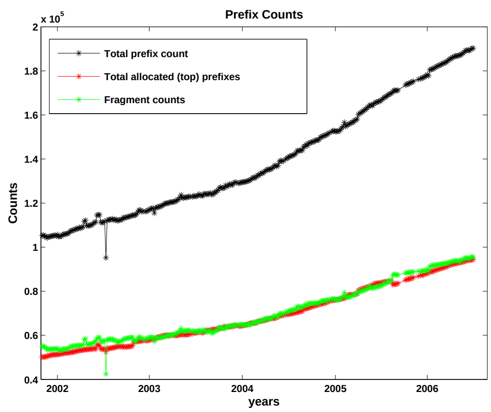

Figure 8: Prefix Counts: Allocated blocks and Fragmented blocks.

This means that the growth in the number of prefix advertisements and BGP table sizes are equally contributed to by address allocations and address fragmentation. The delegation tree continues to grow with continued fragmentation and allocation, but the rate of growth and its impact on the structure and properties of the delegation tree is uniform.

6. Discussion

The ability to study the full delegation of IP addresses is hampered by the length of the address prefixes in the BGP routing table. As previously discussed, addresses blocks with a prefix greater than /24 are usually filtered. This problem is further compounded by the fact that delegations of smaller address blocks often happen within an AS and are hence hidden from public view. For example, an ISP such as AT&T may delegate a portion of its IP address space to a small customer, but they will statically route that block to their routers and not advertise it individually, but rather the advertisement will be implicit within the context of a larger advertised prefix. This may be desirable in that it prevents the organization in question from having to perform their own routing, and such an delegation may be intentionally non-public by companies who wish to keep their arrangements a trade secret; however, it hampers efforts to determine a full delegation structure for the entire address hierarchy. While there is no public source of such data, we might be able to infer some peering and customer-provider or inter-ISP delegations by collecting the AS numbers belonging to a single organization. This may provide futher insight into the delegation structure and may change how the AS graphs correlate. We plan to extend this analysis as a part of our future work.

The fear of exhaustion of IPv4 address space was a prime rationale for the creation of IPv6, which allows for 128-bit addressing. We can project the understanding of IPv4 delegation structure to predict a similar centralized delegation hierarchy for IPv6 with a small fraction of prefixes involved in the delegations. In IPv6, RIRs and large ISPs are allococated with huge /32 and /48 chunks [1]. Further delegations will be in allocations of /48 and /64 length prefixes. Since IPv6 is new, we would observe a initial high allocation rate due to allocations by RIRs and top ISPs. Current data on IPv6 is insufficient to provide a further understanding of the evolution and fragmentation of delegation structure. However, we can expect that fragmentation within these delegated blocks will be larger because of the huge amount of delegated space, but if delegation patterns continue in the same manner as they have, the rate of delegation will stay constant, though lower than what is currently seen in the IPv4 address space given the vastness of the IPv6 range.

As we showed in our analysis, most of the delegation occurring in the current address space is done by a small number of entities, such as national and regional registries or Tier-1 ISPs. For proponents of a public-key infrastructure (PKI) within the Internet, this is a particularly beneficial result. A global PKI is essential for many proposals to secure BGP and address prefixes from theft and hijacking to properly function such as S-BGP [18] and soBGP [21], but one of the primary arguments against such a deployment was that the degree of delegation and the presumed vast number of organizations that would be required to perform delegations would make the deployment intractable. This work demonstrates that because of the centralized nature of address delegation, such fears may be unfounded. A relatively small number of organizations would need to issue certificates, with ICANN/IANA acting as the root for all delegations. Seo et al. [27] consider a PKI that can be used by S-BGP; this work shows that such a scenario is feasible because of the limited number of parties involved, making the entire infrasturcture ultimately manageable.

7. Conclusion

In this paper, we expanded the notion of prefix delegation relationships presented in [20]. We constructed the delegation hierarchy of the IPv4 address space examined the relationship between the Internet topology and the delegation structure, using the Route Views data corpus and processing over 1.6 billion routing table entries. We noted that prefix delegations are largely independent of the Internet topology. We studied the evolution of delegation structure and found it to be fairly stable with a uniform pattern. During the study period, we noticed that the prefix delegations are highly centralized and a large fraction of the address prefixes are allocations by RIRs and top ISPs. We observed that the increase in the number of advertised address prefixes were due to increase in allocations as well as increase in fragmentation. We noticed that the rate of growth of allocations and fragments are similar implying that the rate of fragmentation of the whole address space itself has remained constant. Over the years, address space usage has evolved with fragmentation of existing allocations and addition of new allocations at around the same rate.

By gaining an understanding of the delegation structures and rates of allocation, we can better understand issues of address space exahaustion and fragmentation. With the potentially imminent deployment of IPv6 within the US government and current adoption in Japan, a global deployment of IPv6 may come to be realized within a time period of five years. The address space usage can be extrapolated to grow at a rate that will not exhaust the IPv4 space before this time, and because of the centralized nature of delegations, the potential for deployments of secure routing proposals is possible.

References

[1] ARIN: IPv6 Registeration Services. http://www.arin.net/registration/ipv6.

[2] CAIDA, the Cooperative Association for Internet Data Analysis. http://www.caida.org.

[3] Internet Assigned Numbers Authority. http://www.iana.org.

[4] Internet Protocol V4 Address Space. http://www.iana.org/assignments/ipv4-address-space.

[5] A. Broido and E. Nemeth and K. Claffy. Internet Expansion, Refinement and Churn. In *European Transactions on Telecommunications,13, No. 1, 33-51*, 2002.

[6] A. Broido and K. Claffy. Analysis of Route Views BGP Data: Policy Atoms. In *Proceedings of the Network-Related Data Management (NRDM) workshop*, 2001.

[7] A. Broido and K. Claffy. Complexity of Global Routing Policies. In *Proceedings of the IMA Special Workshop: Mathematical Opportunities in Large–Scale Network Dynamics*, 2001.

[8] A. Broido and K. Claffy. Internet Topology: Connectivity of IP Graphs. In *Proceedings of the SPIE ITCOM*, 2001.

[9] B. Z. R. Liu and D. Massey and L. Zhang. Collecting the Internet AS–Level Topology. In *Proceedings of the ACM SIGCOMM Computer Communications Review*, pages 53–61, New York, NY, USA, 2005. ACM Press.

[10] D. Meyer. The RouteViews Project, May 2003. http://www.routeviews.org.

[11] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. In *Proceedings of ACM SIGCOMM*, pages 251–262, 1999.

[12] G. Huston. IPv4 Address Utilization, 2005. http://www.potaroo.net/papers/2005-03-ipv4.pdf.

[13] G. Huston. Numerology, Nov. 2005. http://www.potaroo.net/ispcol/2005-11/numerology.html.

[14] G. Huston. The BGP Report for 2005, June 2006. http://www.potaroo.net/ispcol/2006-06/bgpupds.html.

[15] L. Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Trans Netw.*, 9(6):733–745, Dec. 2001.

[16] H. Chang and R. Govindan and S. Jamin and S. J. Shenker and W. Willinger. Towards Capturing Representative AS Level Internet Topologies. In *ACM SIGMETRICS*, 2002.

[17] H. Narayan and R. Govindan and G. Varghese. The Impact of Address Allocation and Routing on the Structure and Implementation of Routing Tables. In *Proceedings of the ACM SIGCOMM*, 2003.

[18] S. Kent, C. Lynn, and K. Seo. Secure Border Gateway Protocol (S-BGP). *IEEE Journal on Selected Areas in Communications*, 18(4), Apr. 2000.

[19] L. Subramanian and S. Agarwal and J. Rexford and Randy H. Katz. Characterizing the Internet Hierarchy from Multiple Vantage Points. In *IEEE INFOCOM*, 2002.

[20] P. McDaniel, W. Aiello, K. Butler, and J. Ioannidis. Origin Authentication in Interdomain Routing. *Computer Networks*, 50(16):2953–2980, Nov. 2006.

[21] J. Ng. Extensions to BGP to support secure origin BGP (soBGP). Internet Draft, Oct. 2002.

[22] Q. Chen and H. Chang and R. Govindan and S. Jamin and S. Shenker and W. Willinger. The Origin of Power Laws in Internet Topologies Revisited. In *IEEE INFOCOM*, 2002.

[23] S. Bellovin and R. Bush and T. Griffin and J. Rexford. Slowing routing table growth by filtering based on address allocation policies. June 2001. www.cs.princeton.edu/ jrex.

[24] S. Jaiswal and A. L. Rosenberg and D. Towsley. Comparing the Structure of Power-Law Graphs and the Internet AS Graph. In *ICNP'04*, 2004.

[25] S. Y. Qiu and P. McDaniel and F. Monrose and A. D. Rubin. Characterizing Address Use Structure and Stability of Origin Advertisements in Inter–Domain Routing. In *Proceedings of the 11th IEEE Symposium on Computers and Communications (ISCC)*, 2006.

[26] S. Zhou and R. Mondragon. Analyzing and Modelling the AS Level Internet Topology. In *HET-NETs'03*, 2003.

[27] K. Seo, C. Lynn, and S. Kent. Public-Key Infrastructure for the Secure Border Gateway Protocol (S-BGP). In *IEEE DARPA Information Survivability Conference and Exposition II*, Anaheim, CA, USA, June 2001.

[28] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP topologies with Rocketfuel. *IEEE/ACM Trans. Netw.*, 12(1):2–16, 2004.

[29] X. Meng and Z. Xu and B. Zhang and G. Huston and S. Lu and L. Zhang. IPv4 Address Allocation and the BGP Routing Table Evolution. In *Proceedings of the ACM SIGCOMM Computer Communication Review*, pages 71–80, New York, NY, USA, 2005. ACM Press.

[30] Y. Rekhter and T. Li. RFC 1771 – A Border Gateway Protocol 4 (BGP–4), 1995. http://www.faqs.org/rfcs/rfc1771.html.