Quantifying the Extent of IPv6 Deployment

Elliott Karpilovsky[1], Alexandre Gerber[2], Dan Pei[2], Jennifer Rexford[1], Aman Shaikh[2]

[1] Princeton University
[2] AT&T Labs – Research

Abstract. Our understanding of IPv6 deployment is surprisingly limited. In fact, it is not even clear *how* we should quantify IPv6 deployment. In this paper, we collect and analyze a variety of data to characterize the penetration of IPv6. We show that each analysis leads to somewhat different conclusions. For example: registry data shows IPv6 address allocations are growing rapidly, yet BGP table dumps indicate many addresses are either never announced or announced long after allocation; Netflow records from a tier-1 ISP show growth in native IPv6 traffic, but deeper analysis reveals most of the traffic is DNS queries and ICMP packets; a more detailed inspection of tunneled IPv6 traffic uncovers many packets exchanged between IPv4-speaking hosts (e.g., to traverse NAT boxes). Overall, our study suggests that from our vantage points, current IPv6 deployment appears somewhat experimental, and that the growth of IPv6 allocations, routing announcements, and traffic volume probably indicate more operators and users preparing themselves for the transition to IPv6.

1 Introduction

IPv4, the current Internet protocol, is showing its age. Addresses are becoming scarce, with estimates of exhaustion within the next several years [1]. People are looking toward IPv6, with its 2^{128} possible addresses, as the solution. While there has been pressure to deploy IPv6, NAT technologies have extended IPv4's life. Given the lack of urgency to upgrade, coupled with the administrative and financial overhead of becoming IPv6-enabled, it is difficult to say whether we have moved any closer to a day when IPv6 is so dominant that IPv4 can be "switched off."

Not only has IPv6 deployment has been slower than expected, but our understanding of it is surprisingly limited as well. Questions such as, "are organizations actually using IPv6," "what IPv6 transitional technologies are being used," and, "what applications are using IPv6" remain largely unanswered.

To answer these questions, we looked at a variety of data sources, ranging from regional Internet registry (RIR) logs to BGP dumps to Netflow records to packet header traces. Along the way, we found several "gotchas," where the surface level analysis implies a different conclusion from the in-depth analysis. For example:

- RIR data indicates that IPv6 prefixes are being allocated at near exponential rates, implying strong growth. However, longitudinal BGP analysis shows that nearly half of these allocated prefixes are never announced, and the remainder take an average of 173 days to appear in the global routing system. In other words, many people are acting as IPv6 speculators but not deployers.

– Native IPv6 traffic analysis of the enterprise customers of a US tier-1 ISP shows considerable volume, yet most of the traffic is generated by DNS and ICMP; this indicates a dearth of real IPv6 applications.
– A reasonable amount of tunneled IPv6 traffic is already observed on a US broadband ISP (0.001% of total traffic). However, further analysis indicates that much traffic is between IPv4 clients, implying that IPv6 tunneling technologies are primarily used to circumvent NAT and firewall restrictions.

The paper is organized as follows. Section 2 examines IPv6 address block allocation and compares it against an analysis of address block announcements. Section 3 looks at both native and tunneled IPv6, analyzing the types of technologies used to enable IPv6 communication and the application mix. We discuss related work in Section 4, and conclude with Section 5.

2 Allocation and Announcement of IPv6 Address Blocks

RIR and BGP data are important for understanding how IPv6 addresses are allocated and announced, respectively.

2.1 Data Sources

RIR allocations are important because they indicate the number of institutions requesting blocks, as well the sizes being allocated. For our analysis of IPv6 allocations, we used the ARIN public FTP repository [2] that maintains information about the five regional registries responsible for allocating IPv6 address blocks: APNIC, RIPE, ARIN, LACNIC, and AFRINIC. Date ranges for the different repositories are: 1999-8-13 to 2008-9-25 (APNIC); 1999-8-12 to 2008-9-26 (RIPE); 1999-8-03 to 2008-9-23 (ARIN); 2003-1-10 to 2008-9-22 (LACNIC); and 2004-12-14 to 2008-9-23 (AFRINIC).

In order to analyze how address blocks are announced, we used the RouteViews BGP data archives [3]. We collected routing table (RIB) snapshots at approximately 6 hour intervals from this web site. The BGP data obtained from RouteViews starts on 2003-5-3 and ends on 2008-9-28.

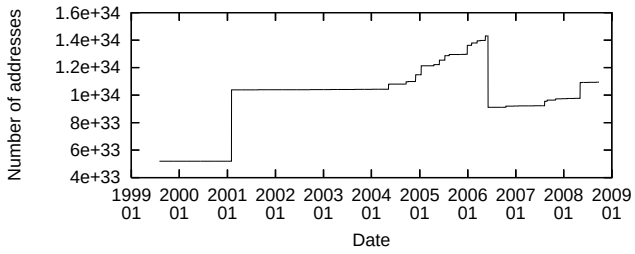

Fig. 1. Number of IPv6 addresses allocated by RIRs. Each prefix is completely de-aggregated.

2.2 Why Address Allocation Statistics are Misleading

Looking at the distribution of allocated prefixes, along with the total number of addresses allocated, seems like a reasonable method for quantifying the extent of IPv6 deployment; however, we find that using such information can be misleading.

First, one can incorrectly conclude that *IPv6 address allocations are very volatile*, as seen by the gigantic spike and dip in the curve. Figure 1 shows the total number of IPv6 addresses assigned by the RIRs. We count the number of allocated addresses by de-aggregating all prefixes into their corresponding sets of IPv6 addresses, and unioning these sets together. A couple points clearly stand out. On 2001-2-1, the number of addresses doubles, due to the allocation of `2002:/16` to the 6to4 transitional technology [4], which reserved a part of the IPv6 space for 6to4; since this is a reservation, it cannot be considered a true measure of IPv6 growth. Likewise, a gigantic drop is seen on 2006-6-6, due to the decommissioning of the 6Bone (`3FFE::/16`). Since the 6Bone was experimental and decommissioned once IPv6 growth was observed [5], it cannot be considered evidence of significant IPv6 constriction.

Second, one can incorrectly conclude that *IPv6 growth has plateaued*. The number of allocated addresses only grew by 20% from 2006-6-7 to 2008-9-26, nearly a 2.5 year period. However, growth is masked by a few extremely large prefixes, that hide the allocation of smaller ones. As of 2008-9-28, of the 2774 allocated prefixes, 31 had a prefix length of 22 or shorter, compared to a median prefix length of 32. Since such large address blocks are allocations (*i.e.*, delegating responsibility of address assignment to local or national registries) as opposed to assignments, they are not true measures of growth. This delegation is the explanation for the plateau, and it is incorrect to draw conclusions about IPv6 growth based on it.

2.3 Drawing Correct Conclusions from the RIR Allocation

What information can be gleaned from the RIR data? After further analysis, we find that statistics concerning the number of prefix allocations provide insight into the deployment of IPv6.

Why Analyzing Prefix Allocations is Better Since looking at the number of allocated *addresses* is not particularly insightful, why would looking at the number of allocated *prefixes* be appropriate? Moreover, is it really fair to look at prefixes independent of their size, lumping the /64s with the /48s and /32s, treating them as all "equal"?

Table 1 shows the distribution of various prefix lengths as a function of year. As time passes, more and more /32 address blocks are present, and the statistics indicate that it is the favorite for recent allocations. In fact, as of 2008-9-28, more than 67% of all prefixes allocated were /32. We believe this heavy bias is due to RIR policy changes in 2004, which made obtaining /32s easier [6–8].

Since so many prefixes are the same size, analyzing allocated prefixes will be roughly fair (since most address blocks are the same size), and the results will not be skewed by the few "heavy hitters" seen before.

Unfortunately, the sub-allocations are not recorded in the RIRs. Thus, if a /32 is allocated to an organization, and that organization re-allocates it to others, only the first

year	allocations	mean	mode	1st quartile	median	3rd quartile
2005-3-3	1183	33.65	32	32	32	34
2006-3-3	1421	33.39	32	32	32	34
2007-3-3	1720	34.19	32	32	32	34
2008-3-3	2179	34.65	32	32	32	34

Table 1. Distribution of prefixes over time. Numbers are cumulative over time.

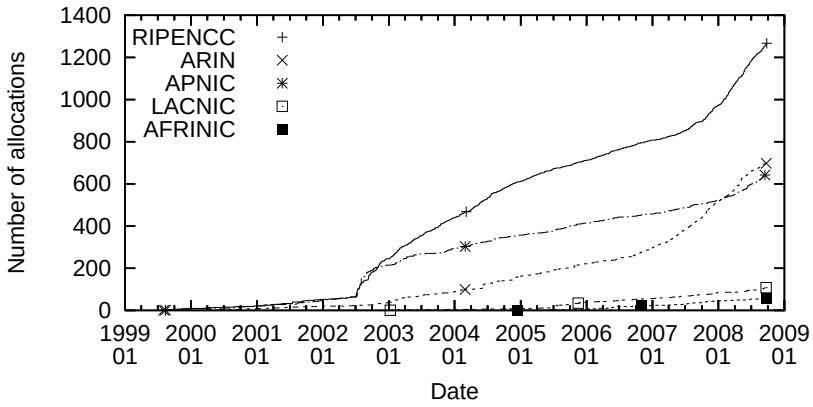

Fig. 2. Growth of the individual registries.

entry is recorded. Thus, our prefix allocation analysis can potentially underestimate the growth of IPv6.

Prefix Allocations Reveal Growth Trends Figure 2 shows allocations of address blocks for the different registries. RIPE is clearly dominant, and shows extremely large growth for 2008. Likewise, ARIN allocations are also increasing at a very fast rate, causing it to surpass APNIC. APNIC has many allocations, but has been growing slowly, and only starts to show signs of acceleration toward the end of the study. LACNIC allocations remain approximately linear. While AFRINIC has very few allocations, it shows signs of acceleration. Cumulatively, there have been nearly 2800 address block allocations. Overall, it would appear as if IPv6 growth has been somewhat stagnant, increasing at a mostly linear rate, until recently.

One point on the graph requires explanation. In July of 2002, RIPENCC and APNIC experienced abnormal growth. Investigation revealed that on July 1st, 2002, RIPENCC and APNIC both instituted policy changes regarding IPv6 allocation [9]; this policy set new guidelines for how to allocate space, as well as guidelines for allocating to new organizations desiring IPv6 space. For example, it defined clear criteria for requesting IPv6 space, as well as a method for organizations to request additional allocations. As such, we believe that these policy changes are responsible for the sudden surge in allocations.

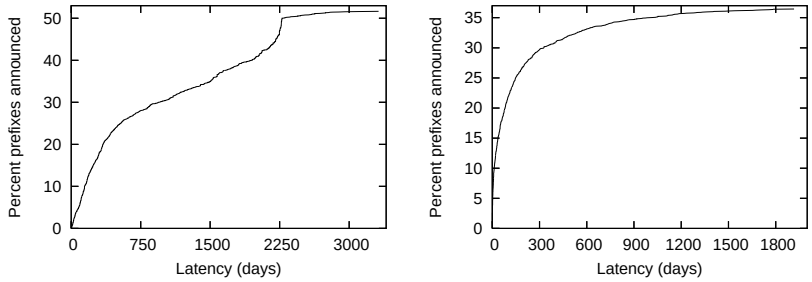

Fig. 3. CDFs of latency for prefixes where BGP data existed. Left graph represents the 52% of prefixes that were never announced. Right graph represents the 36% that were announced. The remaining 12% were allocated before our BGP data begins.

To summarize, IPv6 allocation has only recently started taking off; previous years had mostly linear growth, while current growth could possibly be exponential. However, since we are only at the beginning of the "knee" in the curve, we should be careful in extrapolating too far into the future.

However, address allocations do not imply address usage. Is this allocation really a good measure of IPv6 deployment?

2.4 Prefixes are Allocated, but Often not Used

We turn to BGP routing data, which documents which IPv6 addresses are currently routable on the IPv6 Internet. Figure 3 shows how long it takes institutions to announce their address blocks (or any sub-block) after they've been allocated. We call this "usage latency," and is defined as the difference in time between allocation and the prefix's first appearance in the BGP routing tables.

There are a few points of note about Figure 3. *52% of all allocated prefixes never appear*. Measuring latency for these prefixes is impossible, since it is unknown when they will be used (if ever). Instead, we measure minimum latency, *i.e.*, the minimum amount of time before usage. We find that the average minimum latency is 957 days (which is an underestimate for the true latency!). The kink in the curve corresponds to the policy changes of APNIC and RIPENCC in July of 2002, as mentioned earlier; many addresses allocated during that surge were never used.

When computing true latency for the remaining prefixes, we run across a snag. Approximately 12% of all prefixes were allocated before 2003, when our BGP data begins. As such, it is impossible to accurately measure latency for these prefixes. We ignore such prefixes since they are a small minority.

For the remaining 36% of prefixes, latency averages 173 days. For comparison, we look to a study concerning usage latency in IPv4 in 2005 [10]. The study found that 87% of all prefixes were eventually announced in BGP, and the average latency for these prefixes was 52 days. Thus, there are fewer IPv6 users than their IPv4 counterpart, and they are also much slower to deploy.

Overall, there was some slight variation in latency between regions, but all regions were within a factor of 1.5 of each other. RIPE averaged 141 day latency. LACNIC's latency was 159 days, and AFRINIC's was 177 days. APNIC and ARIN were nearly identical, at 202 and 211 days, respectively.

3 Traffic Analysis in a US Tier-1 ISP

While the RIR and BGP data capture the rate of IPv6 adoption, this ignores three other aspects of IPv6 deployment – how people are transitioning their IPv4 networks to IPv6, how IPv6 addresses are actually assigned to individual machines, and what IPv6 applications are being used. Identifying transitional technologies helps us understand how IPv4 networks connect to the IPv6 world; The upper 64-bits of an IPv6 address identify such mechanisms, as they each have different IP ranges. Observing addresses tells us how organizations assign IPv6 addresses to individual interfaces. Since the low order 64 bits are reserved entirely for host machines, we can use this to see how individual organizations number their devices. Third, to analyze the application mix, we look at the signature *(source port, destination port, protocol)* to map it to an application name.

3.1 Data Sources

To analyze native IPv6 traffic, we use Netflow records collected from an IPv6 Internet gateway router in a US tier-1 ISP with 11 IPv6 BGP neighbors. These records were collected from 2008-4-1 to 2008-9-26, and are taken from the business customers. To analyze tunneled traffic, we collected packet header traces from 2008-7-02 to 2008-8-31 at an access router servicing approximately 20,000 DSL subscribers (different from the business customers) in an ISP. In particular, we analyzed the IPv6 headers within Teredo tunnels. Teredo [11] is an IPv6 tunneling technology created by Microsoft to enable IPv6 communications for Windows users. Due to the prevalence of Windows among typical Internet users, we assume that most tunneled IPv6 traffic destined for these subscribers use Teredo.

Unfortunately, our records are not perfect, and have some holes. Note that 5th, 6th, 9th, 10th, and 11th of July are not analyzed for Netflow. Also note that for the tunneled traffic, data from 2008-7-10 to 2008-7-21, along with 2008-8-19, 2008-8-23, and 2008-8-26 are not included.

3.2 Identifying Transitional Technologies and Address Enumeration

We identify transitional technologies as follows. Teredo uses `2001:0000:` for the first 32 bits [11] of an IPv6 address, making it easily identifiable. 6to4, another popular encapsulating scheme, begin with `2002:`. Although other transitional schemes exist and can be identified (*e.g.*, ISATAP, automatic tunnels, *etc.*), they are quite rare in practice; as such, we lump them together as under the label "other".

To discover how organizations assign addresses to devices, we use the same methodology as presented in [12]. The types of enumeration are: Teredo (Teredo encodes options and routing information in the lower 64-bits), MAC address based (also called

(a) Native IPv6 Records

name	2008-4	2008-5	2008-6	2008-7	2008-8	2008-9
Native IPv6	85.5%	90.5%	87.0%	74.2%	63.5%	75.2%
6to4	12.7%	7.1%	10.6%	23.4%	32.4%	20.8%
Teredo	1.7%	2.3%	2.4%	2.3%	4.1%	4.0%
Other	0.1%	0.1%	0.0%	0.0%	0.0%	0.0%

(b) Tunneled IPv6 Headers

name	2008-7	2008-8
Native IPv6	70.2%	44.2%
6to4	2.8%	4.8%
Teredo	26.7%	50.3%
Other	0.3%	0.7%

Table 2. Monthly averages of different IPv6 technologies seen.

auto-configuration), low (only using the last 12 bits), wordy (using words that can be spelled in hexadecimal, like BEEF), privacy (all bits are randomly set, according to the IPv6 privacy specification [13]), v4 based (when a corresponding IPv4 address influences the choice of host address), and unidentified (for all others).

3.3 Transitional Technologies

The results of analyzing the IP address structure are presented in Table 2. Most of the native IPv6 addresses of the tier-1 ISP tended to communicate with other native IPv6 addresses; approximately 80% of addresses fell into this category. 6to4 addresses were also significant, representing approximately 18% of addresses seen. Teredo addresses constituted approximately 2%, and the remaining technologies were almost negligible. These results also match those found for an analysis done in 2007 [12]. As an important note, the data sets used in our analysis are quite different from those in [12] (which included web server traffic, name server traffic, and traceroutes). Since we have a different vantage point and a different time-frame, yet have the same results, we believe that the technologies used by organizations remain unchanged for the past year.

From the tunneled perspective, we see that Teredo and native addresses are popular. Moreover, around 2008-8, a surge of Teredo-to-Teredo connections is seen.

3.4 Assigning Addresses to Machines

In addition to looking at transitional technologies, we looked at the breakdown of IPv6 address assignment schemes. Table 3 demonstrates the ratios of various host configurations. A few interesting trends emerge. First, IPv4 based addresses decline sharply (although there is a spike in August that remains unexplained). Moreover, privacy extensions remain relatively unused, occupying a small percentage of all addresses (possibly because some operating systems do not enable privacy extensions by default).

3.5 Application Mix

Looking at the application breakdown yielded interesting results, as seen in Table 4. Expected traffic, like web and mail, was surprisingly low – usually between 1% to 8% for web and 1% and 2% for mail. We performed DNS reverse lookups on the few IPv6 addresses that used web protocols and found that popular sites include an IPv6 deployment and tunnel broker and a backbone network for universities. On average,

name	2008-4	2008-5	2008-6	2008-7	2008-8	2008-9
IPv4 Based	49.5%	28.7%	19.3%	5.9%	20.2%	6.1%
Low	22.0%	29.9%	32.5%	36.5%	31.0%	34.8%
Auto-configured	18.6%	29.2%	33.5%	40.3%	31.2%	42.6%
Teredo	1.7%	2.3%	2.4%	2.3%	4.1%	4.0%
Wordy	0.2%	0.2%	0.4%	0.3%	0.8%	0.3%
Privacy	1.0%	1.0%	1.3%	1.7%	1.5%	1.5%
Other	7.0%	8.6%	10.5%	11.8%	11.2%	10.7%

Table 3. Monthly averages of assignment schemes seen for the native IPv6 records.

about 85% of traffic is DNS queries and 8% ICMP messages. Overall, these results are quite surprising. We believe there are two possible reasons. One could be that people are mainly using probing applications over their IPv6 networks, and not actual applications. Another is that operating systems like Windows Vista will send an extra DNS request when IPv6 capabilities are turned on: one requesting the IPv4 address and one requesting the IPv6 address [14]. Thus, the IPv6 interface may send and receive DNS queries but not traffic. Despite the potential inflation of DNS records in our data, there is still very little "real" traffic seen for IPv6. We believe that this demonstrates, for at least this tier-1 ISP, customers view IPv6 as experimental.

For Teredo tunneled traffic, application breakdown was also interesting. Table 4 shows that almost all traffic is unidentifiable UDP or TCP, indicating random port numbers. Given the vast quantity of unidentifiable traffic, and the rise of Teredo pairs, it is likely that these are P2P applications communicating with each other (as random port numbers are characteristic of P2P traffic). Indeed, some applications have turned to Teredo to solve the issue faced by end hosts that are limited by their NAT/firewall technologies when they try to initiate communications with each other; using the Teredo protocol, a client contacts a Teredo server, which acts as a broker agent between Teredo clients, aiding in NAT/firewall hole punching, as well as providing unique IPv6 addresses. Several P2P clients have implemented IPv6 support [15], such as uTorrent and Vuze (formerly Azureus); moreover, uTorrent has the ability to set up Teredo automatically [16]. To summarize, it appears as if considerable tunneled IPv6 traffic is a by-product of applications (such as P2P file-sharing) using Teredo as a mechanism to bypass local NATs and firewalls, simplifying the application developers' jobs.

(a) Native IPv6 Records

name	2008-4	2008-5	2008-6	2008-7	2008-8	2008-9
DNS/Domain	75.5%	86.0%	87.9%	85.3%	88.8%	93.1%
ICMP	11.0%	10.2%	6.9%	6.5%	7.3%	5.2%
Web	8.3%	1.9%	1.3%	2.7%	0.8%	0.4%
Mail	1.3%	0.4%	1.0%	1.5%	0.4%	0.3%
Other	3.9%	1.5%	2.9%	4.1%	2.7%	1.0%

(b) Tunneled IPv6 Headers

name	2008-7	2008-8
Random TCP	30.6%	94.7%
Random UDP	67.3%	3.2%
Web	0.2%	0.03%
Other	1.9%	2.07%

Table 4. Monthly averages of applications; percentages based on number of bytes.

4 Related Work

IPv6 topology has been investigated by CAIDA's scamper work [17], as well Hoerdt and Magoni's Network Cartographer [18]. Because we did not investigate this aspect of IPv6 deployment, we consider our work to be complementary to these studies.

Anomalous BGP behavior has been analyzed through Huston's automatically generated IPv6 reports [19]. These reports include information about routing instability, prefix aggregation, table sizes, and allocation sizes.

Testing the readiness of IPv6-enabled software occurred in February of 2008, when NANOG shut off IPv4 access from their meeting for one hour [20]. It resulted in a severe restriction of services, with end users often needing to re-configure their machines. It revealed that IPv6-enabling software is still somewhat user unfriendly [21]. We believe this work on *how an individual can use IPv6* to be complementary to our work on *how organizations are using IPv6*.

Regarding traffic analysis, Arbor Networks [22] found that IPv6 traffic is growing at the same rate as IPv4 traffic. Savola [23] analyzed 6to4 traffic and found much was experimental, and also noted a rise in P2P applications. Hei and Yamazaki [24] analyzed 6to4 traffic on a relay in Japan and found that TCP traffic dominated UDP, with a considerable amount of HTTP traffic (40% of total). Our work complements these studies because we analyze different data sources, and offer a new perspective by analyzing traffic from a tier-1 ISP.

Finally, David Malone's work on IPv6 addresses analyzed transitional technologies and the assignment of IPv6 addresses to machines [12]. He looked at the breakdown of types of IPv6 addresses (Teredo, 6to4, *etc.*), as well as the classification of the host part of IPv6 addresses. While we do repeat some of the same analysis (and use some of the same techniques), we believe there are key differences between our study and his. We cover broader ground by looking at more data sources: RIR allocations, BGP data, Netflow records, and packet header traces. We also perform additional analysis, such as address space allocation and latency.

5 Conclusion

While IPv6 is beginning to see larger deployments, it still has some significant barriers to overcome. IPv6 is still viewed as experimental by some, and often is deployed in counter-intuitive ways. By analyzing RIR and BGP data, it appears that many allocations are speculatory, and that autonomous systems wait significant amounts of time before actual announcement. Moreover, although IPv6 traffic is growing, our data from a US tier-1 ISP indicated that much of it is still DNS and ICMP packets, indicating a lack of true IPv6 applications from our vantage point; additionally, tunneled traffic analysis shows much of the communication is between IPv4 pairs, implying that applications like P2P file sharing are dominant.

Further work would include a longer study of these characteristics, as well as a topological study involving more end hosts. Moreover, it would be interesting to track operating system developments and their support for various transitional schemes, as well as native support, to better understand how this software shapes the future of IPv6.

References

1. "IPv4 Address Report." http://www.potaroo.net/tools/ipv4/index.html.
2. "ARIN public stats." ftp://ftp.arin.net/pub/stats/.
3. D. Meyer, "University of Oregon Route Views Archive Project." http://archive.routeviews.org/.
4. "IPv6 Global Unicast Address Assignments." http://www.iana.org/assignments/ipv6-unicast-address-assignments, May 2008.
5. B. Hinden, "6bone Phaseout Planning." http://go6.net/ipv6%2D6bone/ngtrans/IETF-56-SanFrancisco/6bone-phaseout.pdf, March 2003.
6. "ARIN Number Resource Policy Manual." http://www.arin.net/policy/nrpm.html, August 2008.
7. "IPv6 Address Allocation and Assignment Policy." http://www.apnic.net/policy/ipv6-address-policy.html, August 2008.
8. "Policy - AfriNIC IPv6 Address Allocation and Assignment Policy." http://www.afrinic.net/docs/policies/afpol-v6200407-000.htm, March 2004.
9. B. Roseman, "ASO Statement Regarding IPv6 Policy Adoption." http://www.icann.org/en/aso/ipv6-statement-11jul02.htm, July 2002.
10. X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, and L. Zhang, "IPv4 address allocation and the BGP routing table evolution," *SIGCOMM Computer Communication Review*, 2005.
11. "Teredo Overview." http://technet.microsoft.com/en-us/library/bb457011.aspx, January 2007.
12. D. Malone, "Observations of IPv6 Addresses," in *Passive and Active Network Measurement*, Springer Berlin / Heidelberg, 2008.
13. T. Narten, R. Draves, and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6." http://tools.ietf.org/html/rfc4941, September 2007.
14. "Domain Name System Client Behavior in Windows Vista." http://technet.microsoft.com/en-us/library/bb727035.aspx, September 2006.
15. "SixXS - IPv6 Deployment & Tunnel Broker :: IPv6 BitTorrent Clients." http://www.sixxs.net/tools/tracker/clients/, September 2008.
16. "forum.utorrent.com / uTorrent 1.8 released." http://forum.utorrent.com/viewtopic.php?id=44003, August 2008.
17. B. Huffaker and K. Claffy, "Caida : research : topology : as_core_network." http://www.caida.org/research/topology/as_core_network/ipv6.xml.
18. M. Hoerdt and D. Magoni, "Distribution of multicast tree states over the IPv6 network topology," in *2004 IEEE Conference on Communications*.
19. G. Huston, "IPv6 Reports." http://bgp.potaroo.net/index-v6.html.
20. P. Smith, "IPv6 Hour at NANOG42." http://www.nanog.org/mtg-0802/ipv6hour.html, January 2008.
21. J. Doyle, "IPv6 Hour at NANOG: A Follow-Up." http://www.networkworld.com/community/node/25276, February 2008.
22. S. Iekel-Johnson, C. Labovitz, D. McPherson, and H. Ringberg, "Tracking the IPv6 Migration." http://www.arbornetworks.com/IPv6research, 2008.
23. P. Savola, "Observations of IPv6 Traffic on a 6to4 Relay," in *ACM SIGCOMM Computer Communication Review*, January 2005.
24. Y. Hei and K. Yamazaki, "Traffic analysis and worldwide operation of open 6to4 relays for IPv6 deployment," in *2004 Symposium on Applications and the Internet*.