



Tracking the IPv6 Migration

GLOBAL INSIGHTS FROM THE LARGEST STUDY TO DATE ON IPv6 TRAFFIC
ON THE INTERNET

About this Report

This Arbor Networks® research report presents the results of an extensive one-year study designed to measure the usage of Internet Protocol version 6 (IPv6) on the Internet. Billed as the “next-generation” Internet protocol, IPv6 is intended to replace the 20-year-old IPv4 protocol. The migration to IPv6 is critical given that IPv4 is rapidly running out of available addresses—a fact that could eventually trigger the collapse of the Internet. Previously, only anecdotal evidence and limited reports were available to track this migration. This research report presents, for the first time, a truly global perspective on the amount of IPv6 traffic on the Internet. Developed by Arbor Networks in partnership with more than 90 network services and content providers from around the world, the report is believed to be the largest and most comprehensive study of IPv6 traffic to date.

About Arbor Networks

Arbor Networks is a leading provider of secure service control solutions for global networks. Its customers include over 70 percent of the world’s ISPs and many large enterprises. Arbor solutions deliver best-in-class network security and visibility, along with the power to improve profitability by deploying differentiated, revenue-generating services. By employing flow-based and deep packet inspection (DPI) technologies, Arbor solutions measure and protect the entire network—from the network core to the broadband edge.

Arbor solutions include the Arbor Peakflow® family, the de facto standard for flow-based network security and analysis, and the Arbor Ellacoya eSeries platforms for broadband service optimization, ranked among the highest performance DPI systems. The combination of these solutions enables Arbor to detect and mitigate network-based attacks across the entire infrastructure, without the exorbitant cost of deploying inline devices pervasively throughout the network. Arbor Peakflow products are typically deployed in the core for comprehensive threat detection and mitigation. Arbor Ellacoya eSeries products, which are deployed at the broadband edge of wireline and wireless networks, utilize DPI and application classification—enabling providers to offer application- and subscriber-aware services.

Introduction: Why IPv6?

Internet Protocol version 6—or IPv6 for short—is intended to be the next primary communications protocol for packet-switched inter-networks. As such, it is the designated successor to IPv4. The transition to IPv6 is considered important primarily because current projections indicate that Internet Assigned Numbers Authority (IANA) and the various Regional Internet Registries (RIRs) will run out of available IPv4 addresses within the next few years.¹ IPv6 has over 28 orders of magnitude more potential addresses than IPv4, which should be sufficient for quite some time.

The IPv4-to-IPv6 transition has received a great deal of media attention over the past year, partly due to the recent U.S. Office of Management and Budget (OMB) deadline for IPv6 adoption,² and partly due to the more obvious impending exhaustion of the IPv4 address pool. At its peak this past June, print media around the world published nearly 3,000 articles a month on IPv6 (almost twice the number of IPv4 articles).³

To date, studies of IPv6 have been limited in scope. Previous reports have looked at deployment challenges facing IPv6 adoption⁴ and related traffic metrics. These metrics include the number of ASNs with IPv6 BGP announcements, the fraction of the Alexa Top 500 Web sites using IPv6, and IPv6 DNS queries as a percentage of IPv4 DNS load. One study collected traffic statistics from select IPv6-IPv4 tunnel endpoints.⁵ All of these metrics are indirect estimators of what we actually want to know: How much IPv6 traffic is on the Internet?

Arbor Research Study: Parameters and Methodology

To answer this question, Arbor Networks conducted a one-year study that is the first to offer a global and longitudinal perspective on the amount of IPv6 traffic on the Internet today. Our measurements included anonymized traffic data from more than 90 ISPs representing a broad cross section of global Tier 1 IP network service providers and regional Tier 2 ISPs, as well as large content providers, hosting companies and broadband access providers. In all, our data collection covered 2,393 peering and backbone routers, along with 278,797 customer and peering interfaces. By the end of the year, we were seeing 4 Tbps of inter-domain Internet traffic. We believe this is the largest study of IPv6 traffic to date. Table 1 provides a breakdown of reporting service providers per geographic region.

| Region | Reporting Service Providers |
|--------------|-----------------------------|
| Americas | 65 |
| EMEA | 27 |
| Asia Pacific | 6 |

Table 1: Regional Breakdown of Reporting Service Providers

Source: Arbor Networks, Inc.

Traffic data was collected via flow export from peering, aggregation and customer-facing routers at each participating network by commercial flow measurement appliances,⁶ and by in-line or off-ramped deep packet inspection (DPI) devices deployed at key points in the network. Traffic measurements were made in 5-minute intervals by averaging the number of bytes and packets seen from all flows and packets during each interval. Measurements were later aggregated into daily averages by a round-robin database for viewing year-long trends.

Our research looked at all IPv6 traffic—both native and tunneled, including Teredo, which encapsulates IPv4 traffic in UDP datagrams using UDP port 3544. We found a peak of only 12 Mbps of Teredo traffic, representing around 10 percent of the IP protocol 41 traffic.

¹ <http://www.potaroo.net/tools/ipv4/index.html>

² E. Castelli, "Agencies ready for next-generation Internet, OMB reports," *Federal Times*. July 1, 2008. <http://federaltimes.com/index.php?S=3607993>

³ Based on a MetaNews query for print media articles containing "IPv6" and "IPv4" in the first 30 words.

⁴ M. Tatipamula, P. Grossetete, H. Esaki, "IPv6 integration and coexistence strategies for next-generation networks." *IEEE Communications Magazine*. January 2004.

⁵ Hei, Y. Yamazaki, K., "Traffic analysis and worldwide operation of open 6to4 relays for IPv6 deployment." Applications and the Internet, 2004. Proceedings. 2004 International Symposium. August 2004

⁶ Arbor Networks Peakflow Products, <http://www.arbornetworks.com>

One challenge to measuring the amount of IPv6 traffic on the Internet today is that many routers are still not capable of exporting flow records for native IPv6 traffic. This is generally due to either a lack of support in the router's software or hardware, or to router configuration issues. In the case of many routers, for example, native IPv6 or IPFIX flow export requires the use of NetFlow v9, but many ISPs still use NetFlow v5 instead due to either hardware, software, licensing or configuration issues. Our measurements showed a negligible amount of native IPv6 traffic. However, since we do not currently have a good estimate for the number of monitored routers in our study that are capable of exporting flow records for native IPv6 traffic, we are unable to conclude whether this is due to a true lack of native IPv6 traffic or simply due to the limitations of the network reporting infrastructure itself.

By excluding native IPv6 traffic, it is possible that we are undercounting the total amount of IPv6 traffic reported by the participating networks. There are no comprehensive measurements of how much of today's IPv6 traffic is tunneled as opposed to native. At least one recent study has shown that 90 percent of IPv6 traffic on the Internet today is tunneled over IPv4,⁷ while others have estimated that tunneled traffic represents around 25 percent of total IPv6 traffic.⁸ Regardless of the actual percentage, it seems that tunneled IPv6 traffic is a significant fraction of all IPv6 traffic. Knowing the amount of tunneled IPv6 traffic also provides a likely upper bound on the total amount of IPv6 traffic, whether native or tunneled. We therefore believe that our measurements provide the best available estimate of overall IPv6 traffic on the Internet to date.

IPv6 Traffic Measurements

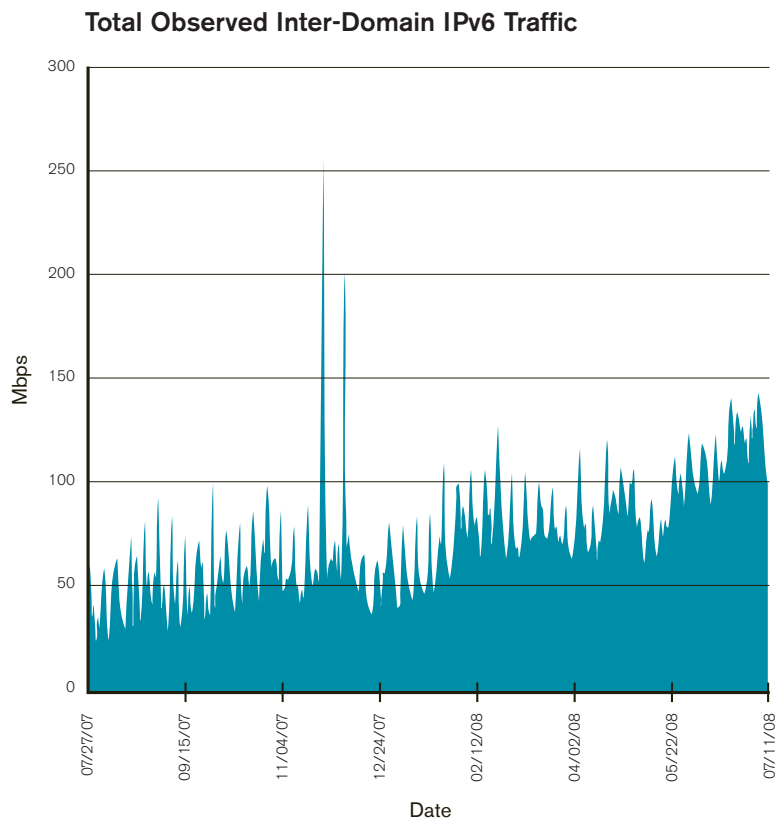


Figure 1: Total Observed Inter-Domain IPv6 Traffic

Source: Arbor Networks, Inc.

⁷ J. Palet, A. Vives, "Measuring Real Global IPv6 Traffic." RIPE 55. October 2007. <http://www.ripe.net/ripe/meetings/ripe-55/presentations/palet-v6.pdf>

⁸ G. Huston, G. Michaelson, "Measuring IPv6 Deployment." RIPE 56. May 2008. http://www.ripe.net/ripe/meetings/ripe-56/presentations/Huston-Measuring_IPv6_Deployment.pdf

Figure 1 (page 3) shows the total amount of inter-domain IPv6 traffic that is tunneled over IPv4 as a function of time. The numbers are provided as an aggregate across all monitors. The good news is that we observed steady growth in the volume of IPv6 traffic across our one-year trace: from roughly 50 Mbps during the early fall of 2007 to approaching a peak of 150 Mbps during the summer of 2008. More precisely, there was a 113 percent increase in the mean inter-domain IPv6 traffic observed during the first quarter of our study and the final quarter of our study.

If we assume the lowest reported percentage of tunneled versus native IPv6 traffic cited above (25 percent), this still provides an estimated upper bound of all IPv6 traffic of 600 Mbps by the summer of 2008, compared to a total of 4 Tbps of all IPv4 inter-domain traffic.

Interestingly, there are a couple of obvious spikes around the fall of 2007. The second spike may correspond to IETF 70 (December 2-7, 2007) and some IPv6 deployment exercises that were conducted there.⁹ We have not yet identified a source for the first spike.

Note that this data has not been normalized with regard to the number of monitored routers in order to compensate for additional coverage across existing participants. As a result, some of the traffic increase may be due to new participating networks or new routers within existing networks being added to the anonymous data reporting systems. However, we have also identified significant increases of IPv6 traffic within individual ISPs over the same period. Therefore, we believe that the data does indicate an overall increase in IPv6 traffic in the Internet over the course of the year.

“We have worked with Arbor Networks on a number of initiatives and they offer ISPs the opportunity to come together in a way that fosters collaboration as well as a better understanding of industry challenges. IPv6 represents the future of the Internet, and adoption of this protocol will continue to increase. We look forward to working with the ISP community and our customers to ensure a smooth transition.”

John Hayduk, Chief Technology Officer,
Tata Communications




Figure 2: Fraction of Inter-Domain Traffic that is IPv6

Source: Arbor Networks, Inc.

⁹ <http://asert.arbornetworks.com/2008/04/ipv6-eat-your-own-dogfood-report/>

While Figure 1 (page 3) indicates that the raw amount of tunneled inter-domain IPv6 traffic has grown steadily over the past year, so has the total amount of Internet traffic. Figure 2 (page 4) shows tunneled IPv6 traffic as a fraction of overall IPv4 Internet traffic. Over our one-year trace, tunneled IPv6 traffic represented only 0.0026 percent of overall IPv4 traffic. Moreover, there was a 10 percent drop in the mean fraction during the first quarter of our study and the final quarter of our study. In other words, the fraction of IPv6 traffic as a function of aggregate traffic on the Internet has remained stagnant over the past year, despite significant press coverage and government mandates.

While this 10 percent decrease may represent an apparent decline in IPv6 tunneled traffic levels in comparison to aggregate traffic rates, it may also indicate an increase in native IPv6 traffic on the Internet.

Conclusions

The deployment of IPv6 is important to the future health of the Internet. The dwindling supply of available IPv4 addresses has been broadly documented, as have the challenges faced by those wishing to migrate to IPv6. In terms of hard data, however, industry and academia have had little visibility into the rate of IPv6 deployment in the Internet. Previous studies have been limited to indirect estimators of the key question: How much IPv6 traffic exists on the Internet today, and is it growing?

This Arbor Networks research report presents the largest study of IPv6 traffic on the Internet to date. Our monitors were distributed across 87 different network services and content provider networks, covering over 2,000 routers and nearly 300,000 interfaces. Near the conclusion of our one-year measurement period, we were seeing nearly 4 Tbps of traffic when averaged over a full day. We therefore believe that our measurements provide the best available estimate of overall IPv6 traffic on the Internet.

Based on our measurements, we have observed that:

- The amount of aggregate inter-domain IPv6 Internet traffic appears to be increasing.
- IPv6 traffic is still a tiny percentage of overall Internet traffic. By one conservative estimate, there were about 600 Mbps of inter-domain IPv6 traffic by the end of July 2008, compared to 4 Tbps of IPv4 traffic.
- The proportion of IPv6 versus IPv4 traffic stayed roughly the same over the past year.

This last point is particularly interesting. Our data seems to indicate that there is currently no significant migration of users from IPv4 to IPv6. However, since overall IPv4 traffic grew significantly, it also implies that IPv6 usage is growing at roughly the same rate as IPv4 (or perhaps slightly slower). As with IPv4, this growth in IPv6 traffic may stem from new users migrating to IPv6 and existing users increasing their IPv6 utilization to meet their demand for high bandwidth applications such as streaming video.

Going forward, Arbor intends to perform a more detailed analysis of native IPv6 traffic on the Internet, as well as analyzing both native and tunneled intra-domain traffic. We also plan to look at the application distribution of native IPv6 traffic and geographic distributions of sources and destinations to better understand how the use of IPv6 may differ from IPv4.

For more industry breaking news and trends, please visit the Arbor Security Blog at <http://asert.arbornetworks.com>

About the Authors

Dr. Scott Iekel-Johnson, Principal Software Engineer, Arbor Networks

scottij@arbor.net

Scott Iekel-Johnson has extensive experience in computer networking and distributed systems. Since joining Arbor Networks in 2001, he has worked on designing and implementing the Peakflow SP product line. Prior to joining Arbor, Scott was at the University of Michigan, where he conducted research into building highly scalable, fault-tolerant distributed systems. He has also participated in research at IBM's T.J. Watson Research Laboratory and with the IPMA project at Merit Network, Inc.

Scott holds a Ph.D. and MSE in Computer Science from the University of Michigan, and a BSE in Electrical Engineering and Computer Science from Duke University.

Dr. Craig Labovitz, Chief Scientist, Arbor Networks

labovitz@arbor.net

Craig Labovitz brings extensive experience in network engineering and research to Arbor Networks. Before joining Arbor, he served as a network researcher and scientist for the Microsoft Corporation. Previously, he spent nine years with Merit Network, Inc. and the University of Michigan as a senior backbone engineer and director of the Research and Emerging Technologies group. His work at Merit included design and engineering on the NSFNet backbone and Routing Arbiter projects. He also served as the director of several multimillion dollar grants from the National Science Foundation for network architecture and routing protocol research.

Dr. Labovitz received his PhD and MSE from the University of Michigan.

Danny McPherson, Vice President and Chief Security Officer, Arbor Networks

danny@arbor.net

With over 15 years experience in the Internet network operations, security and telecommunications industries, Danny McPherson brings extensive technical leadership to Arbor Networks. Today he is a main contributor to the company's overall strategy and product architecture. Prior to joining Arbor, he was with Amber Networks. Previously he held network operations and architecture positions for nearly a decade at internetMCI, Genuity (acquired by GTE Internetworking), Qwest Communications and the US Army.

McPherson has been an active participant in Internet standardization since 1996. Currently he is a member of the Internet Architecture Board (IAB) and co-chairs the IETF's PWE3 WG. He also serves on the MPLScon Advisory Board and the FCC's Network Reliability and Interoperability Council (NRIC), and is quite active in the network and security operations and research communities.

McPherson has authored a significant number of books, Internet protocol standards, network and security research papers and other documents related to Internet routing protocols, network security, Internet addressing and network operations.

Haakon Ringberg, Research Assistant, Arbor Networks

haakon@princeton.edu

Haakon Ringberg is a fifth-year Ph.D. student at Princeton. He has collaborated with great researchers both in academia and at the industry labs of Thomson Paris, AT&T Research and Arbor Networks. His research interests span networked systems with a focus on network security.

Haakon expects to receive his Ph.D. from Princeton in 2009. He holds an M.A. from Princeton in computer science and a B.A. from Cornell University with majors in computer science and philosophy.



Corporate Headquarters

430 Bedford Street
Lexington, Massachusetts 02420
Toll Free (USA) +1 866 212 7267
T +1 781 684 0900
F +1 781 768 3299

Europe

T +44 208 622 3108

Asia Pacific

T +65 3627 7152

www.arbornetworks.com

Copyright ©1999-2008 Arbor Networks, Inc.
All rights reserved. Arbor Networks and the
Arbor Networks logo are trademarks of Arbor
Networks, Inc. All other brands may be the
trademarks of their respective owners.

SR/IPV6/US/0808

About Arbor Networks

Arbor Networks is a leading provider of secure service control solutions for global business networks, including more than 70 percent of the world's Internet service providers and many of the largest enterprise networks in use today. Arbor is addressing the most strategic issues for service providers—security and service control; delivering best-in-class network protection and the means for delivering revenue-generating, differentiated secure services and service plans. Arbor allows service providers to employ both flow-based and DPI-based technologies to enable measurement and protection of the entire network, from the core to the broadband edge. Arbor also maintains the world's first globally scoped threat analysis network—ATLAS—which uses technology embedded in the world's largest ISP networks to sense and report on comprehensive worldwide threat intelligence.