

IPv6 User Experience: An Assessment of Top U.S. University Domains

Dr. John Pickard, East Carolina University

Dr. Pickard is an Assistant Professor at East Carolina University in the College of Engineering and Technology. He teaches undergraduate and graduate Information and Computer Technology (ICT) courses within the Department of Technology Systems. Dr. Pickard plays an active role in building positive and sustainable industry relationship between the college, local businesses, and industry partners. Current industry recognized certifications include; Cisco Certified Network Professional, Microsoft Certificated Professional, EMC Information Storage and Management, IPv6 Forum Certified Engineer (Gold), IPv6 Forum Certified Trainer (Gold), and Cisco Certified Academy Instructor. Dr. Pickard received his Ph.D. in Technology Management at Indiana State University. He also holds an MBA from Wayland Baptist University and a B.S. in Professional Aeronautics from Embry-Riddle University. Research interests include: IPv6, IPv6 adoption, wireless sensor networks, and industry-academia partnerships.

Mr. Dale Drummond, East Carolina University

Dale Drummond is an Undergraduate Student at East Carolina University in the College of Engineering and Technology. He is currently pursuing his Bachelor of Science in Information and Computer Technology with a concentration in Computer Networking. Mr. Drummond is currently Team Lead/Sr. Network Management Engineer at MCNC, the company which operates the North Carolina Research and Education Network (NCREN). Research interests include IPv6, IPv6 adoption, Application Performance Monitoring, DevOps, and Network Operations Optimization.

Dr. Philip J. Lunsford II, East Carolina University

Phil Lunsford received a B.S. in Electrical Engineering and a M.S. in Electrical Engineering from Georgia Institute of Technology and a Ph.D. in Electrical Engineering from North Carolina State University. He is a registered professional engineer and is currently an Associate Professor at East Carolina University. His research interests include cyber security, network performance, and the cross-discipline application of technologies.

Ciprian Popoviciu

Dr. Ciprian Popoviciu, MBA, has over 20 years of experience working in various technical and leadership roles in the networking and cloud computing area. Prior to starting Nephos6 he managed the architecture team of Cisco's Engineering Infrastructure Services organization where he defined the strategy and led the execution of the internal DC consolidation and transition to cloud. He is an industry recognized domain expert in IPv6 who, over the past 15 years worked on the strategy, architecture, implementation and validation of IPv6 adoption by large service providers and enterprises Worldwide. Ciprian worked with various governments on defining national level and agency level IPv6 strategies and standards. Ciprian authored two Cisco Press books and co-authored 5 IETF RFCs and multiple patents. Ciprian is an IPv6 Forum Fellow and a Cisco Champion. He is a frequent speaker to industry events and conferences and he is an active contributor to various technology and scientific publications.

IPv6 User Experience: An Assessment of Top US University Domains

Introduction

Internet Protocol version 4 (IPv4) has served as the dominant Internet protocol for more than thirty-five years. However, its ability to support and sustain the growth of the Internet is no longer in question. In fact, the last remnants of the global IPv4 address pools are already drying up. As Table 1 shows, four of the five Regional Internet Registries (RIRs) have exhausted their IPv4 address pools and are operating under exhaustion policies that restrict IPv4 address allocations. Only the RIR for Africa, AFRINIC, has yet to reach full exhaustion. At the time of writing AFRINIC is allocating from its final /8 and has a projected exhaustion date of 09 November 2019 [1].

Table 1. Regional Internet Registry Last /8 Allocation Information.

RIR	Description	Last /8 Assignment	
		Start	Exhaustion
AFRINIC	The Internet Numbers Registry for Africa	04/03/2017	11/09/2019*
APNIC	The Asia-Pacific Network Information Centre	04/19/2011	05/27/2014
ARIN	American Registry for Internet Numbers	01/30/2014	09/24/2015
LACNIC	The Internet Addresses Registry for Latin America & Caribbean	05/19/2014	02/15/2017
RIPE	Regional Internet Registry: Europe, Middle East, Central Asia	09/14/2012	03/17/2018

* projection as of April 2019.

IPv6 is the next generation Internet protocol and is the successor to IPv4. The principal motivation for the development of IPv6 was to address the limited number of unique Internet address supported by IPv4. The 32-bit address field of the IPv4 protocol header allows for only 4.3 billion unique addresses. IPv6, however, boasts a 128-bit address field that supports up to 340 undecillion (340 with 36 zeros) unique IP addresses. Additionally, IPv6 comes with other technical improvements: a simplified header for faster router processing, a stateless auto-configuration mechanism for address provisioning, improved support for mobility, and built-in support for quality-of-service and security [2].

Despite the technical superiority of IPv6, the Internet's migration to the new protocol has been anything but smooth. Many have expressed doubts, and some still do, that IPv6 will reach full adoption and replace IPv4 as the Internet's dominant protocol. However, empirical data suggests that Internet IPv6 adoption has entered a phase of rapid acceleration [3]. A recent study by [4] found that the number of IPv6 users on the Internet has reached the early majority level of adoption and full adoption could occur as early as December 2024. Figure 1 shows the number of users accessing Google over IPv6 reached 26.31% in January of 2019 [5].

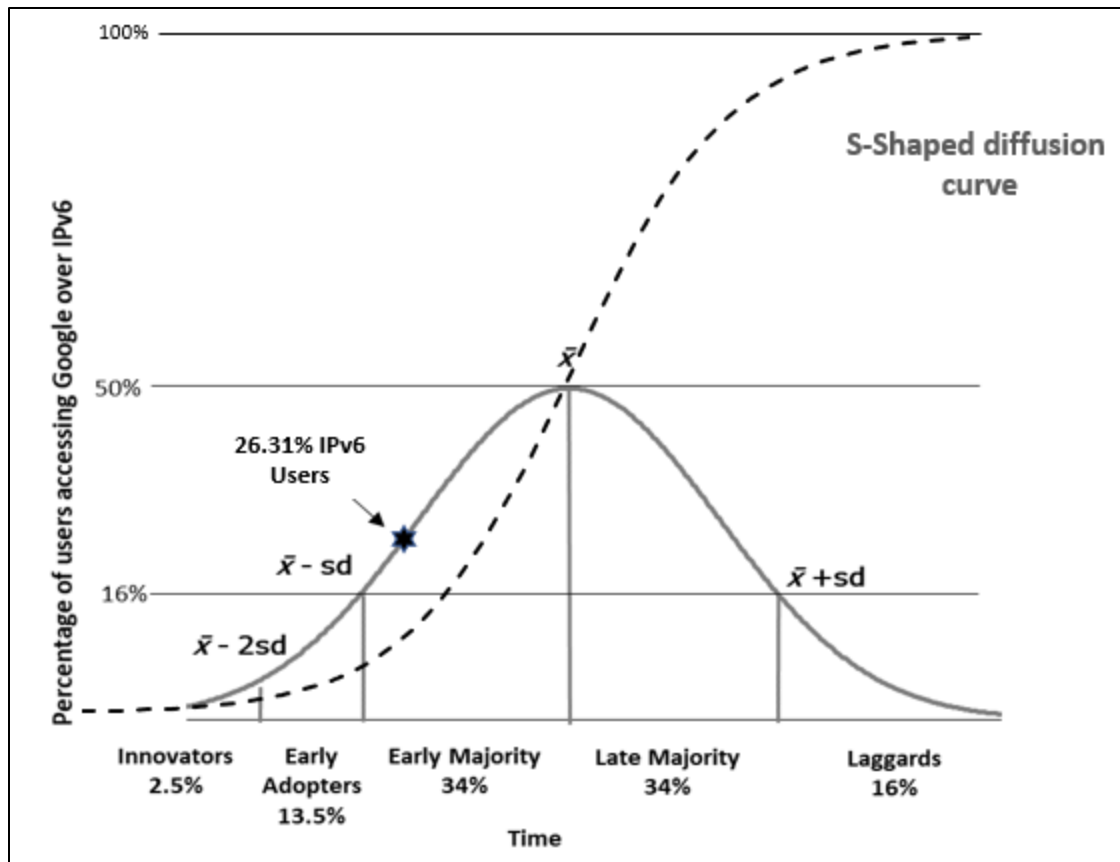


Figure 1. The adopter distribution normal curve partitioned into adopter categories overlaid with the S-shaped diffusion curve [6]. The number of IPv6 users accessing Google over IPv6 in January 2019 reached 26.31% [5].

The number of IPv6 users captured by Google is the result of default operating system behavior and growing support for IPv6 among hardware vendors and service providers. Today, all major desktop and mobile operating systems come with IPv6 enabled and many broadband Internet Service Providers (ISP) are deploying IPv6 unknown to their customers. Enabling IPv6 and IPv4 concurrently in a dual-stack scenario allows users to both automatically take advantage of the browser IPv6 fast start capabilities for sites that support IPv6 and reach the resources that are only IPv4 reachable. This dual-stack setup lets the user access both IPv6 and IPv4 resources without having to know how to configure IPv6 or even be aware that their device is configured to use it. When connecting to a dual-stacked web site, IPv6 enabled hosts will prefer IPv6. If IPv6 cannot establish a TCP session within a defined period (300 ms in some browsers) the host will attempt the session over IPv4.

As the number of users with IPv6 connectivity grows, web site operators can expect increased IPv6 traffic to their dual-stacked sites. While users will not be aware of the protocol used to establish the connection, they will be very aware of poor web site performance and broken links. They will likely lay the blame for the poor user experience on the web operator, quite possibly leading to a negative impact on the organization brand. Poor IPv6 performance can also lead to an organization being listed on Google blacklist [9]. Again, damaging the brand and translating into increased business risk.

Unfortunately, web site operators often lack information on the quality of service provided over IPv6 to users accessing their web content [10]. Most methods of IPv6 monitoring are related to IPv6 adoption and not performance and provide no insight as to how IPv6 is performing. In this paper we provide a quantitative analysis of the availability and accessibility of web services over IPv6. Additionally, we present a methodology for assessing the quality of service experienced by the user accessing the web content of top US universities over IPv6.

Related Works

The adoption and performance of IPv6 on the Internet is an extensively studied topic. Studies typically fall into two categories: empirical studies measuring IPv6 availability and performance and analytical studies that attempt to model factors that aid or hinder IPv6 adoption. In this section we review only a sample of representative empirical studies.

A study by [11] quantified IPv6 adoption and performance by measuring access to content of the Alexa top 1 million websites. The authors found performance over IPv4 and IPv6 was similar when autonomous system (AS) paths were the same. They also found that less efficient AS paths were responsible for most of the instances where IPv6 performance lagged that of IPv4. They were not, however, able to identify a common factor shared by websites with IPv6 performance that outperformed IPv4. The authors [12] also assessed the performance of the Alexa top 1 million websites over both IPv4 and IPv6 and had similar findings. Performance of the two protocols was comparable when the AS-path was the same. However, IPv6 performance tended to be much worse when the path differed. They also noted that the average IPv6 AS-path length was skewed by a single large AS, Hurricane Electric.

Other IPv6 performance studies using the Alexa top websites includes [13] who measured IPv4 and IPv6 connectivity of dual-stacked hosts using TCP connect time. They found that the absence of Content Delivery Network (CDN) caches for IPv6 lead to relatively higher TCP connection times compared to IPv4. [14] compared the similarity of the Alexa top 100 dual-stacked websites when accessed from IPv6 and IPv4. Their findings revealed that 27% of the sites had at least one webpage element that failed over IPv6, and that 9% had more than half of webpage elements that failed over IPv6. In a more recent study [15] used packet delay, packet loss, and packet reordering measurements to analyze the IPv6 performance of dual-stacked websites from the Alexa top 1 million. Like previous studies, they found correlation between IPv4 and IPv6 performance with AS path.

Finally, [16] assessed the user experience accessing the top 100 Alexa top dual-stacked hosts by calculating average HTTP load times over both IPv6 and IPv4. Findings showed that 74% of sites had IPv6 HTTP load times within 10% to that of IPv4, or faster. Also, of interest in their findings was that the physical path of the packets as seen by traceroutes correlated more than AS paths with differences in HTTP load times.

Methodology

A list of 1812 recognized US higher-education institutions was gathered from uniRank [17]. The web domain of each university was queried for the presence of an IPv6 AAAA record in the

global DNS. To accomplish this, the domain name of each institution was loaded into a python script which checked each domain name for an IPv6 AAAA record. The script automated DNS queries for IPv6 AAAA records to two resolvers: Google DNS at 8.8.8.8 and OpenDNS at 208.67.222.222. Both resolvers returned 204 AAAA records. A review of the AAAA records revealed seven duplicates due to some of the institutions having remote campuses that use the same domain, leaving 197 unique university domains with IPv6 AAAA records in the global DNS.

The 197 domains with AAAA records were separated into two lists. Each list of domains was loaded into a network monitoring system, Smartmon.io, and subjected to three tests over both IPv4 and IPv6: ping, http availability, and traceroute from three geographic locations. The domains were split into two lists for testing due to resource limitations of the virtual machines hosting the monitoring agents.

The agents running the Smartmon.io software act like client machines where they request a URL for a dual-stack web server over HTTP performing measurements on the response. This tests the implement of “happy eyeballs” [7] and most closely reflects the user experience when navigating to dual-stack sites. Once the measurements are completed, they are sent to a central server for storage and eventual analysis.

The 197 IPv6 reachable domains were polled at 60-minute intervals for a period of 24 hours from four network monitoring agents deployed in Toronto, New York City, and San Francisco. The use of multiple agent location vantage points helps avoid biases associated with an individual location. The agents at each location were deployed in Virtual Machines (VMs) hosted by Digital Ocean. Each VM ran on CentOS 7.3 with 1 CPU, 512MB of memory, 20GB of storage on a Solid-State Drive (SSD), 1TB of transfer data, and was enabled for both IPv4 and IPv6. Digital Ocean was chosen as the Virtual Machine provider due to the reasonable and deterministic pricing model for their small VMs; their geographic footprint, which allows testing from multiple locations; the robustness in their implementation tools, including an API; and the ability to configure resource monitoring for each Droplet from the Digital Ocean dashboard.

At each 60-minute polling interval, the agents captured the DNS name query time, TCP/IP session establishment time, HTTP load time, and application load time over both IPv4 and IPv6 for each domain. During this process, the agents also recorded the HTTP waterfall, ping results, traceroute data, and the Autonomous System path (AS_PATH) for each domain. Table 2 describes each recorded variable.

Table 2. Description of variables recorded by each network monitor agent [4].

Variable	Description
DNS response time	The amount of time required for DNS responses to be returned to the client. This includes all DNS queries for embedded content as well. DNS latency is an important measure of how a user perceives responsiveness of the DNS server and ultimately how they perceive the speed or performance of accessing web resources [18].
TCP connection time	The amount of time required to establish the transport layer connection including other time intervals such as Secure Socket Layer (SSL) handshakes. The user experiences this as the time it takes for a browser to establish a connection with a web server [19].
HTTP response and load time	Consists of the time to perform the HTTP GET and the time to load the requested page including all images, scripts, and third-party resources [20].
HTTP waterfall	Graphical display showing the roundtrip time between server and browser for each object, including text, images, and JavaScript contained on the target website [21].
Traceroute	A TCP/IP utility that allows a user to trace a network connection from one location to another, recording every hop along the way. When a traceroute is run, it returns a list of network hops and displays the host name and IP address of each connection. It also returns the amount of time it took for each connection to take place, usually in milliseconds [22].
AS_PATH	A well-known BGP path attribute which identifies the autonomous systems through which routing information carried in a BGP UPDATE message has passed [23].

At the end of the 24-hour data collection period, the raw data from the agents was compiled into a .csv file for analysis. The IPv4 load times and IPv6 load times (from all agents to each website) were averaged and domains that experienced 1000 ms longer IPv6 load times than IPv4 were further investigated for AS_PATH, Traceroute, and HTTP waterfall inconsistencies. The purpose was to determine which variable or variables attributed to IPv6 performance lagging behind that of IPv4. The 1000 millisecond time was chosen based on the “two-second rule” which states that the average user abandons a page trying to load after waiting for two seconds [24]. For this paper we focused on those sites whose IPv6 load times were worse than the corresponding IPv4 load times by at least fifty percent of the average two-second abandonment threshold. This allowed us to focus on the difference between the two protocols, rather than eliminating some websites based on overall load time.

Each of the 197 dual-stacked domains were also evaluated for webpage elements that fail to load over IPv6. First, the HTTP waterfall of each university website was reviewed for failing elements. Figure 2 shows an example HTTP waterfall for a university website with failed elements. The waterfall shows two linked elements that fail. The first is the link to the university’s “Give” page. The second is a link the “MyXUniversity” resource login page.

Next, each university website was accessed using a web browser from a PC with only IPv6 access to the Internet. Each menu on the website main page was accessed as well as selected submenus to pages considered to be important to students and parents such as Academics, Admissions, Athletics, and Library. Of the 197 university websites evaluated all 197 had at least one element that failed over IPv6. And nine websites failed to load over IPv6 altogether. It should be noted that in many cases the elements failing over IPv6 were linked the third parties that did not support IPv6.

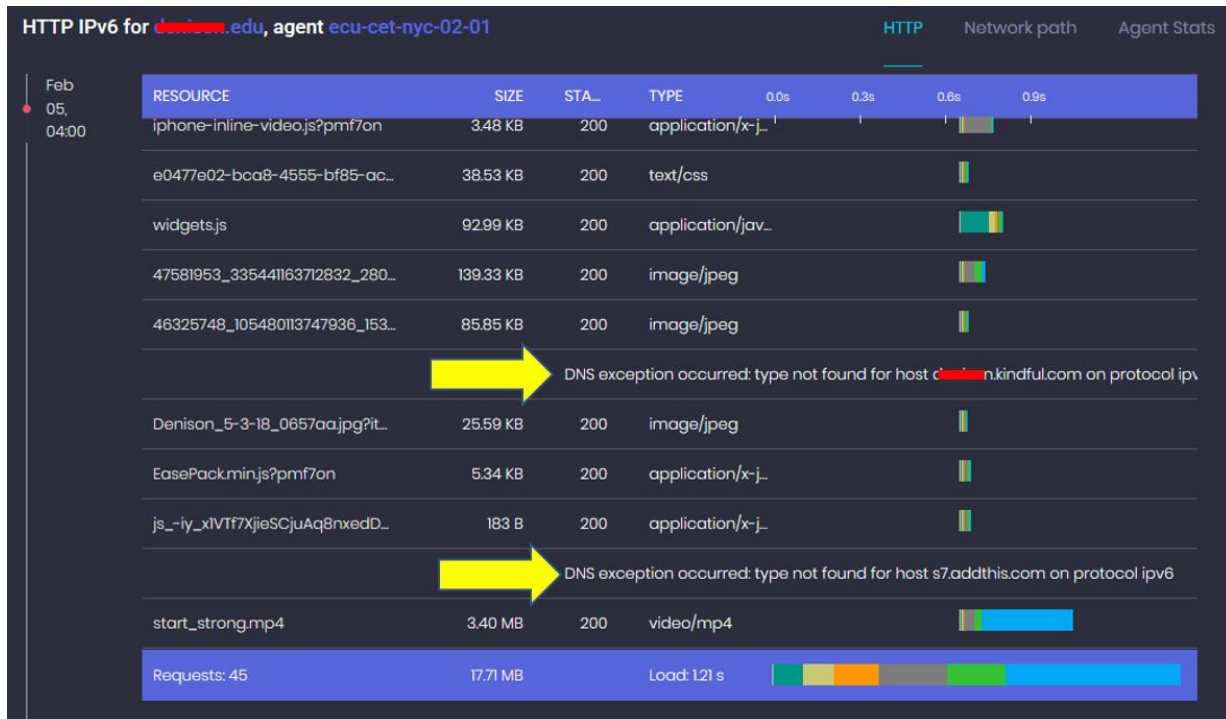


Figure 2. Example TCP waterfall of a university website showing two elements that failed to load over IPv6.

Findings

The first finding is from our quantitative analysis of the 1812 university domains examined. Figure 3 shows that of 1812 university domains only 197 (11%) returned an IPv6 AAAA record. Nine of the 197 domains had websites that failed to load over IPv6. The remaining 188 websites each had at least one element that failed over IPv6. In most instances, the failures were attributed to external links to third party content. For example, many universities use “The Common Application” which is a third-party undergraduate college admission application system that applicants use to apply to member colleges and universities. Common Application does not support IPv6 and therefore this element fails over IPv6 on every university web site that uses it. Other common elements that failed over IPv6 on the web sites of universities were library links, links to athletic pages, and links to philanthropic pages.

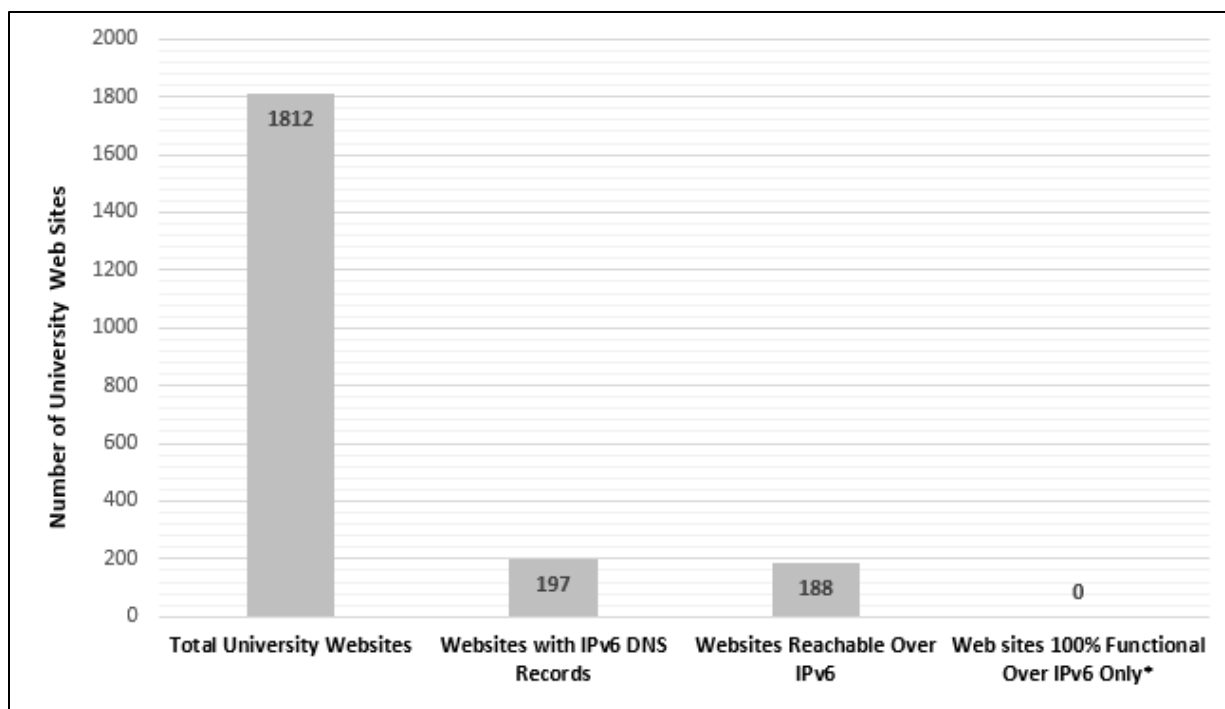


Figure 3. Number of domain websites in each stage of data analysis. Of note is that all domain web sites had at least one element that failed when accessed from and IPv6 only host. *The failure is not attributed to the protocol itself but rather the element did not support IPv6.

The next findings are from a quantitative analysis of HTTP load times. Of the 188 dual-stacked domains accessible over IPv6, 37 (~20%) had HTTP load times faster over IPv6 than over IPv4, three of which were faster by >1000ms. Conversely all other sites had HTTP load times faster over IPv4 than IPv6, with 14 at least 1000ms faster. Figure 4 shows a scatter plot of the average load times of the 188 domain websites. Positive numbers indicate that IPv4 HTTP load times were faster than IPv6. Negative numbers indicate the IPv6 HTTP load times were faster than IPv4.

Comparisons of the HTTP waterfall and traceroute paths for sites that had > 1000ms differences between IPv4 and IPv6 HTTP load times were made. Findings of this investigation revealed that the three sites with significantly faster IPv6 load times (>1000ms) all had elements that failed DNS resolution over IPv6 to hosted content. In each instance these failed elements were associated with content hosted on Amazon Web Services (AWS), Vimeo video content, and Content Delivery Networks (CDNs). Investigations into the 14 sites with significantly faster IPv4 load times (>1000ms) revealed that each had elements with longer HTTP wait times. HTTP wait time is the time a client waits after sending an HTTP request until the first byte is received. No correlation could be found between HTTP load times and traceroute paths.




Figure 4. The millisecond difference between IPv4 and IPv6 HTTP load times for domains. Negative numbers indicate faster IPv6 HTTP load times than IPv4. Positive numbers indicate IPv4 HTTP load times faster than IPv6.

Limitations

Although this study highlights the progress of IPv6 adoption and identifies a method for the qualitative analysis of user experience, it has some limitations associated with the methodology.

The scope of the user experience analysis was limited to US university domains. Analysis could be expanded to include industry and government domains for a more complete picture of the quality of IPv6 enablement on the Internet. Also, user experience monitoring agents were installed on Virtual Machines (VMs) in three North American cities hosted by a single provider, Digital Ocean. Other cloud providers such as Amazon Web Services, Google Cloud Platform, or Microsoft Azure also provide similar services and our assumption is that they would have provided comparable results. In future studies VMs could be hosted on multiple providers and in more geographic locations to ensure there is no single provider or location bias.

Conclusions and Future Work

We have presented a methodology to evaluate user experience accessing web content on the Internet and have applied this methodology to the domains of 1812 top Universities in the US. Our findings revealed that fewer than 11% of the 1812 university domains are advertising availability over IPv6. Deeper analysis showed that all university websites reachable over IPv6 had at least one element that failed over IPv6. Failure and poor performance can cripple the experience for users behind IPv6-only networks and users whose systems choose to use IPv6. While in most cases the failing elements were linked to third-party content that was not IPv6 capable and not operated by the university, the user will likely not make such distinctions.

As a growing number of users connect over IPv6, Organizations, including institutions of higher education, must measure and monitor access to their online content to ensure that it is not only reliably accessible over IPv6 but that it meets user expectations for performance. In short, the performance over IPv6 should be at least as good as IPv4. This can be accomplished through user experience monitoring with end-to-end measurements that consider all aspects of accessing applications and services

We believe that our methodologies provide significant insight into the current state of IPv6 diffusion on the Internet and we intend to continue to measure and systematically monitor the IPv6 diffusion phenomenon as it progresses. Our future work includes extending our investigation into the root causes of the disparate performance behavior experienced access web content over IPv6 and IPv4.

References

- [1] Huston, G. (2019). Retrieved April 23, 2019, from the IPv4 Address Report: <https://ipv4.potaroo.net/>
- [2] Deering, S., & Hinden, R. (2017). Internet protocol, version 6 (IPv6) specification. RFC 8200.
- [3] Nikkha, M., & Gurin, R. (2016). Migrating the internet to IPv6: An exploration of the when and why. *IEEE/ACM Transactions on Networking*, 24(4), 2291-2304.
- [4] Pickard, J., Angolia, M., & Drummond, D. (2019). IPv6 Diffusion Milestones: Assessing the Quality and Quantity of Adoption: An Assessment of Adoption and Quality of Services. *Journal of International Technology and Information Management*. [not yet in print]
- [5] Google IPv6 statistics (n.d.). Retrieved from <https://www.google.com/intl/en/ipv6/statistics.html>
- [6] Rogers, E. (2003). *Diffusion of Innovations* (5th ed.). New York, N.Y.: Simon & Schuster, Inc.
- [7] Wing, D., & Yourtchenko, A. (2012). Happy Eyeballs: Success with Dual-Stack Hosts. RFC 6555.
- [8] Bajpai and Schonwalder (2016). Measuring the effects of happy eyeballs. In *Proceedings of the 2016 Applied Networking Research Workshop* (pp. 38-44). ACM.
- [9] Popoviciu, C., (2016, April 27). Moving towards a better (not just a bigger) Internet. APNIC blog. Retrieved September 2, 2018 from <https://blog.apnic.net/?s=popoviciu>.
- [10] Colitti, L., Gunderson, S., Kline, E., & Refice, T. (2010). Evaluating IPv6 adoption in the Internet. *Passive and Active Measurement* (pp. 141-150). Springer Berlin Heidelberg. doi:10.1007/978-3-642-12334-4_15

- [11] Nikkhah, M., Gurin, R., Lee, Y., & Woundy, R. (2011, December). Assessing IPv6 through web access a measurement study and its findings. In *Proceedings of the Seventh Conference on emerging Networking EXperiments and Technologies* (p. 26). ACM.
- [12] Dhamdhere, A., Luckie, M., Huffaker, B., Elmokashfi, A., & Aben, E. (2012). Measuring the deployment of IPv6: Topology, routing and performance. *Proceedings of the 2012 ACM Internet Measurement Conference*, pp. 537-550.
- [13] Bajpai, V., & Schönwälder, J. (2015, May). IPv4 versus IPv6—who connects faster?. In *IFIP Networking Conference (IFIP Networking)*, 2015 (pp. 1-9). IEEE.
- [14] Eravuchira, J., Bajpai, V., Schonwalder, J., & Crawford, S. (2016). Measuring Web Similarity from Dual-stacked Hosts. In *Network and Service Management (CNSM)*, 2016 12th International Conference on (pp. 181-187). IEEE.
- [15] Li, F., Wang, X., Pan, T., & Yang, J. (2017). A Case Study of IPv6 Network Performance: Packet Delay, Loss, and Reordering. *Mathematical Problems in Engineering*. <https://doi.org/10.1155/2017/3056475>.
- [16] Pickard, J., Southworth, J., & Drummond, D. (2017). The IPv6 Internet: An Assessment of Adoption and Quality of Services. *Journal of International Technology and Information Management*, 26(2), pp. 48-64. <http://scholarworks.lib.csusb.edu/jitim/vol26/iss2/3/>
- [17] UniRank University Rankings (n.d.). Retrieved from <https://www.4icu.org/us/>
- [18] McDonald, D. (2017, June 28). Why you should care about DNS latency. Akamai Blog. Retrieved September 2, 2018 from <https://blogs.akamai.com/2017/06/why-you-should-care-about-dns-latency.html>
- [19] Sexton, P. (2015, August 17). TCP Connection Time. Retrieved September 2, 2018 from <https://varvy.com/performance/tcp-connection-time.html>
- [20] Pingdom (n.d.) Difference between response time and page load time. Retrieved 2 September 2018 from <https://help.pingdom.com/hc/en-us/articles>
- [21] Bixby, J. (2010, April 9). Waterfalls 101: How to understand your website's performance via waterfall chart. Retrieved 2 September 2018 from <http://www.webperformancetoday.com/2010/07/09/waterfalls-101/>.
- [22] Christensson, P. (2006). Traceroute Definition. Retrieved 2018, Sep 5, from <https://techterms.com>
- [23] Rekhter, Y., Ed., Li, T., Ed., & S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>

[24] Galletta, D. F., Henry, R., McCoy, S., & Polak, P. (2004). Web site delays: How tolerant are users?. *Journal of the Association for Information Systems*, 5(1), 1.