

5-1-2019

IPv6 Diffusion Milestones: Assessing the Quantity and Quality of Adoption

John Pickard

East Carolina University, pickardj@ecu.edu

Mark Angolia

East Carolina University, angoliam@ecu.edu

Dale Drummond

MCNC, ddrummond@mcnc.org

Follow this and additional works at: <https://scholarworks.lib.csusb.edu/jitim>



Part of the [Communication Technology and New Media Commons](#), [Digital Communications and Networking Commons](#), [Management Information Systems Commons](#), and the [Technology and Innovation Commons](#)

Recommended Citation

Pickard, John; Angolia, Mark; and Drummond, Dale (2019) "IPv6 Diffusion Milestones: Assessing the Quantity and Quality of Adoption," *Journal of International Technology and Information Management*. Vol. 28 : Iss. 1 , Article 1.

Available at: <https://scholarworks.lib.csusb.edu/jitim/vol28/iss1/1>

This Article is brought to you for free and open access by CSUSB ScholarWorks. It has been accepted for inclusion in *Journal of International Technology and Information Management* by an authorized editor of CSUSB ScholarWorks. For more information, please contact scholarworks@csusb.edu.

IPv6 Diffusion Milestones: Assessing the Quantity and Quality of Adoption

John Pickard
(East Carolina University)

Mark Angolia
(East Carolina University)

Dale Drummond
(MCNC)

ABSTRACT

There are currently two versions of Internet Protocol (IP) in use today, IP version 4 (IPv4) and IP version 6 (IPv6). The original version, IPv4, was standardized in the early 1980s as part of the Defense Advanced Research Project Agency Internet program and became the official Internet protocol in 1983 (Kleinrock, 2010). IPv6 was standardized in 1995 as its successor to provide enhanced capabilities and address IPv4 technological limitations, most notable of which was the anticipated exhaustion of address space (Deering & Hinden, 1995). While the two protocols have some functional similarities, they are distinct and not backward compatible; IPv4-only devices cannot communicate directly with IPv6-only devices and vice-versa. Consequently, organizations wishing to take full advantage of the enhanced features of IPv6 must upgrade their entire network infrastructure and end devices to support IPv6, while at the same time maintaining IPv4 support for legacy systems that will not or cannot be upgraded. The costs and risks associated with upgrading an entire network to support a new protocol with no intrinsic return on investment has acted as a disincentive for IPv6 adoption. To be sure, the transition of the Internet to IPv6 has certainly taken a leisurely pace over the past twenty years. Given the slow pace of adoption, it is understandable that many doubted, and may still doubt that IPv6 will ever become the dominant Internet protocol and replace IPv4. However, in line with diffusion of innovations theory, it is the case with many innovations that potential adopters do not perceive any relative advantage, thus leading to a particularly slow adoption take-up rate. This is especially true with communications technologies that have high interdependence and require a critical mass of users before adoption becomes self-sustaining and rapidly accelerates (Rogers 2003). The goal of this paper is to provide empirical evidence showing that IPv6 adoption has reached critical mass and is now in a phase of accelerating

adoption projected to continue. A methodology for monitoring the quality of IPv6 enablement and global IPv6 support is also provided so that the user experience over IPv6 can be assessed against the IPv4 baseline.

KEYWORDS: IPv6, Internet, Diffusion, IP

INTRODUCTION

The principal technological limitation of IPv4 is a 32-bit address field in the protocol header limiting the number of unique IPv4 addresses to 4.3 billion. This is not enough addressing space to support the growth of next generation information technologies such as the Internet of Things, Software Defined Networks, Cloud computing, and Mobility. The reality is that last remnants of the global IPv4 address pools are drying up and neither an IPv4 only nor a dual-stacked IPv4/IPv6 environment will sustain growth in the long term (Alghatrifi & Khalid, 2018). In February of 2011, the Internet Assigned Numbers Authority (IANA) exhausted its global IPv4 address pool by making a final allocation to each of the five Regional Internet Registries (RIRs) (NRO, 2011). Since 2011, four of five RIRs have likewise exhausted their IPv4 address pools and are now operating under exhaustion policies restricting IPv4 address allocations (Huston, 2018). Table 1 shows the actual IPv4 exhaustion date for each RIR. Note that only the RIR for Africa, AFRINIC, has yet to reach exhaustion. At the time of writing it is currently allocating from its final /8 and has a projected exhaustion date of 16 July 2019 (Huston, 2018).

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	07/16/2019*
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 June 2018.

IPv6, the next generation protocol for the Internet, is the successor to IPv4. It is designed to overcome the addressing limitations of IPv4 and to support continued Internet growth through a greatly expanded address space. IPv6 boasts a 128-bit address field supporting up to 340 undecillion (340 with 36 zeros) unique IP addresses. Other technical improvements are also built into the protocol including 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 (Deering & Hinden, 2017).

Despite technical improvements, years of championing by pundits, (Classe, 2003; Khan & Sindi, 2012; Ladid, 2009; Popoviciu & Dini, 2006), various government mandates (Coleman, 2014; Doyle, 2008; Garretson, 2005; Wu, Wang, & Yang, 2011), and adoption initiatives, ("Internet Society," n.d.; "World IPv6sd Launch," n.d.) IPv6 has yet to replace IPv4 as the dominating Internet protocol. According to Nikkhah and Guerin (2016), IPv6 adoption has gone through a three-phase evolution: 1) stagnation, spanning from 1995 to 2009, 2) emergence, spanning from 2009 to 2011, and 3) the current phase, acceleration, which began in 2011 when IANA announced the exhaustion of its IPv4 address pool. Indeed, the pace of IPv6 adoption is accelerating, however the question remains when or if it will reach a state of full adoption and displace IPv4 as the dominant Internet protocol. This paper aims to offer insight into the answer to this question.

To better understand the slow rate of adoption experienced by IPv6 and predict where it may be in the future, we turn to Rogers' Diffusion of Innovation Theory (Rogers, 2003). It is our hypothesis that the adoption of IPv6 exhibits the characteristics of what Rogers' terms a preventive innovation, as it lowers the probability of a future negative event that may or may not occur, i.e. the shortage or exhaustion of IP addresses. Such innovations tend to have particularly slow uptakes of adoption because potential adopters see no immediate relative advantage to adoption. The reason for this phenomenon, according to Rogers (2003), is that new communications technologies create interdependence among adopters, known as network externalities, and they are of little use unless others also adopt. A good analogy of this phenomenon is the introduction of the telephone. The earliest adopters had a limited number of people to call, and not until mass adoption occurred did the technology become ubiquitous.

The innovation adoption process does not terminate with the decision to adopt a new technology. Rogers (2003) found "empirical evidence supplied by researchers" (p.189), indicating that once an adoption decision is made, adopters enter a confirmation stage where they seek reinforcement for the decision. If the

technology fails to perform as expected, this may lead to dissatisfaction and discontinuance of use or replacement with an alternative technology. Applied to the case of IPv6, achieving mass adoption does not guarantee that the protocol will meet user expectations of reliable access to content and services or that performance will be on par with IPv4. Such an outcome could lead to reduced motivation of Internet stakeholders to provide services over IPv6 and thus prolong the migration of the Internet to IPv6 (Eravuchira, Bajpai, Schonwalder, & Crowford, 2016; Nikkhah & Guerin, 2016).

While the availability of content over IPv6 is within the control of the Internet Content Providers (ICPs), how that content is accessed, i.e., over IPv4 or IPv6, and the quality of the user experience often are not (Popoviciu, 2016). Other Internet stakeholders between the ICP and end user, such as Internet Service Providers (ISPs) and Content Delivery Networks (CDNs), have a significant impact on the user experience. Therefore, a comprehensive measure of IPv6 diffusion on the Internet must measure not only the density of user adoption but also the quality of service experienced by the user (Popoviciu, 2016). Failure to monitor the user experience risks not meeting user expectations and potentially negatively impacting perceived value.

In this paper we extend our previous work (Pickard, Angolia, & Chou, 2018) investigating IPv6 adoption on the Internet. From our research we present three significant findings. First, empirical evidence showing that after two decades of slow uptake, IPv6 adoption has reached the level of critical mass plotted on the normal diffusion curve. Second, we confirm projections of accelerated IPv6 adoption previously published by Pickard et al. (2018) using a predictive model of IPv6 adoption. Finally, we share a methodology for assessing the quality of the user experience when accessing web content, and explore factors that contribute to quality and performance inconsistencies between IPv6 and IPv4.

This paper builds upon the existing body of research combining empirical measurements of IPv6 availability and performance with an analytical re-evaluation of IPv6 adoption against Rogers' innovation diffusion model. This research is valuable and timely to organizations seeking to assess the risks and benefits of migration to IPv6 by providing a methodology for qualitative analysis of the IPv6 enabled infrastructure. Additionally, this paper offers significant contribution through its analysis into the current levels of IPv6 adoption (as of June 2016), its forward-looking insight into the probable progression of IPv6 diffusion on the Internet, and providing ICPs a methodology into accessing the quality of user experience.

RELATED WORKS

A review of literature revealed a significant amount of research offering valuable data on the migration of the Internet to IPv6. Relevant research used for this paper take one of two approaches, empirical studies measuring IPv6 availability or performance and, analytical studies seeking to identify factors affecting IPv6 adoption. Some empirical studies took a narrow approach measuring IPv6 availability and performance with a focus on individual components or stakeholders. For example, Colitti, Gunderson, Kline, and Refice (2010) developed a methodology to measure IPv6 adoption from the perspective of a single web site operator, Google.com. Their findings revealed that IPv6 adoption varied across geographic regions, was increasing rapidly, and that latency over native IPv6 was comparable to that of IPv4 when connecting to IPv6-only and dual-stack hosts. They also found that latency was often negatively impacted when transition mechanisms such as 6to4 and Teredo tunnels were in the data path. Their measurements are still ongoing and are published daily on the Google IPv6 Stats web site ("Google IPv6 Statistics," n.d.).

Another empirical study, by Nikkhah, Guerin, and Woundy (2011), used access to web content as the primary metric to quantify IPv6 adoption and performance on the Internet. They deployed monitoring tools to assess the performance of the Alexa top 1 million websites over both IPv4 and IPv6. The authors found performance was similar over both protocols when the autonomous system (AS) paths were the same. They also found that less efficient AS paths were responsible for instances where IPv6 performance lagged that of IPv4. They were not, however, able to identify any common property shared by sites exhibiting better IPv6 performance. A similar study by Dhamdhere et al. (2012) also found IPv6 performance similar to that of IPv4 if the forwarding AS-path was the same. Their research also showed that IPv6 adoption was higher in the Europe and Asia Pacific regions, and that a single AS, Hurricane Electric, was significantly more prevalent in the IPv6 topology than the most predominant AS in the IPv4 topology, suggesting that the average IPv6 AS-path length may be significantly skewed by a single large AS.

Czyz et al. (2014) took a broad approach to measuring IPv6 adoption and performance using 12 metrics. The authors found that IPv6 adoption indicators, although all showing increasing adoption, varied in the rate of increase across global regions and across all 12 metrics. From these findings they concluded that a broad approach, observing differences across multiple metrics, is essential to fully understanding the true state of IPv6 adoption. Li, Wang, Pan and Yang (2017) used metrics of packet delay, packet loss, and packet reordering to analyze the IPv6

performance of dual-stacked websites from each of the five RIRs. Their findings revealed that packet delay and loss was similar over IPv6 and IPv4 when the AS-level Path was the same. When packet delay and loss were notably higher over IPv6, the performance and the number of ASes in the path was found responsible.

Determining the status of IPv6 deployment and identifying factors inhibiting deployment were the goals of Domingues, Friacas, and Veiga (2007). The authors assessed deployment levels by examining four metrics: (1) ASes found in the BGP forward information base, (2) RIR prefix allocations, (3) Internet core peering, and (4) IPv6-enabled top-level domains. Their investigation revealed incongruence in the networks and number of ASes and network links seen in the IPv6 and IPv4 Internet cores. They found only 2.60% of ASes announced in the IPv6 core as compared to the IPv4 core. They also found uneven distribution of IPv6 ASes between global regions with 51% from the European region, and 48% split evenly between the North American and Asia Pacific regions, with less than 1% seen from the African and Latin American regions.

In a more recent study, Nikkhah and Gurin (2016) assessed the progress of IPv6 adoption across Internet service providers, Internet technology developers, Internet content providers, and end users. They found IPv6 adoption followed a three phase progression of stagnation, emergence, and acceleration. They concluded that low initial demand for IPv6 enabled products along with the lower quality of those products, as compared to those enabled for IPv4, appeared largely responsible for initial reluctance on the part of ISPs and ICPs to adopt IPv6 – which in turn deterred users and prolonged the Internet's migration to IPv6.

Rogers' diffusion of innovation theory was applied by Hovav and Schuff (2005) and Dell et al. (2007) in an effort to identify the drivers and barriers to IPv6 adoption. Both studies similarly concluded that IPv6 adoption decisions are mostly influenced by perceptions of the usefulness of its features and environmental conduciveness. Thirteen years later, Wang and Zander (2018) extended the work of Hovav and Schuff by examining the effects of organizational factors on the Internet Standards Adoption (ISA) model when applied to IPv6 adoption in Australia and China. They found that the organizational factors of complexity and top management support affected IPv6 adoption decisions in both countries, but that normative pressure had more influence in China than in Australia.


Dell (2010) applied the economic theories of exhaustible resources and permit markets to explain the slow progress of IPv6 adoption. Dell concluded that significant IPv6 diffusion would only occur after the exhaustion of the IPv4 address pools. Subsequent studies published after the exhaustion of the IANA global IPv4

address pool by Czyz et al. (2014), Beeharry and Nowbutsing (2016), and Pickard et al. (2017) provide evidence confirming this conclusion.

THEORETICAL FRAMEWORK

The theoretical framework that follows is reproduced from our previous work, (Pickard et al., 2018) and presented here for background. Research by Everett Rogers (2003) led to the establishment of a foundational theory on innovation diffusion through social systems. A key finding of his study was that organizations go through an innovation-decision process like that of individuals. The process begins with gaining initial knowledge of an innovation, then proceeds through forming an attitude about the innovation, deciding to adopt or reject the innovation, implementing the innovation, and finally confirming the decision. Potential adopters move through the process based on their level of innovativeness and do not all adopt an innovation at the same time. Thus, the adoption of innovation usually follows a normal bell-shaped distribution curve plotted over time as shown in Figure 1.

Figure 1. The adopter distribution normal curve partitioned into Rogers' five adopter categories (Rogers, 2003).

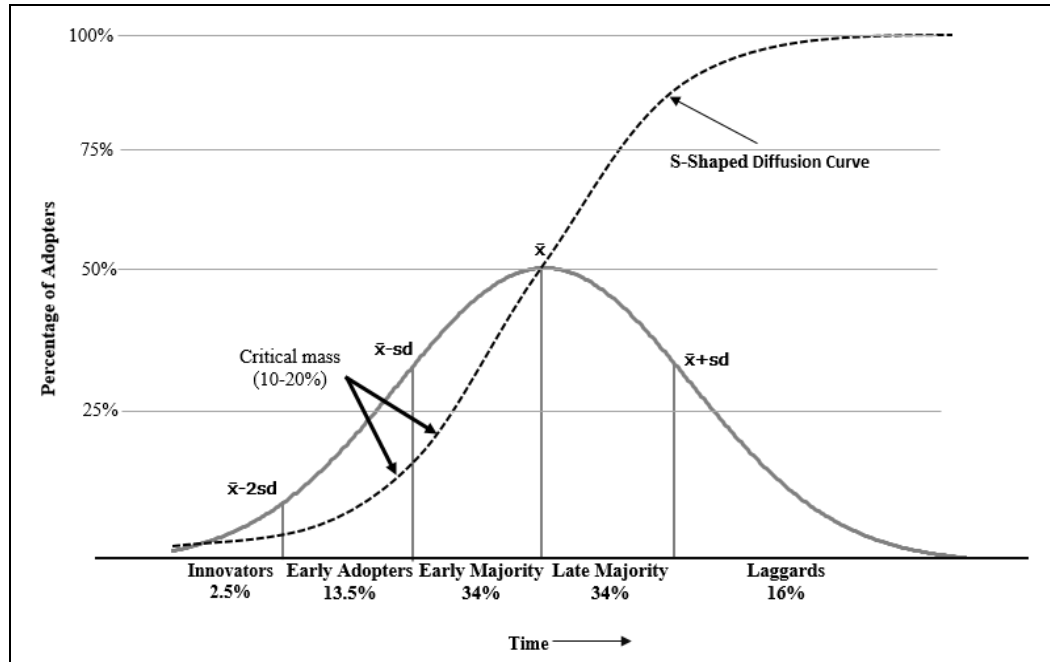


Rogers (2003) used the mean (\bar{x}) and the standard deviation (sd) to map adopter classifications onto the normal distribution curve to define a population's percentage of classification and associated thresholds. The mean rate of adoption

was established at a 50% marketplace share defined by a new technology's adoption. Vertical lines under the curve mark off the standard deviations on either side of the mean into Rogers' five adopter categories. The first 2.5% of users to adopt new technology are classified as "*innovators*" and occupy the extreme left tail of the normal curve, starting at a zero (introduction) point and extending approximately to minus two standard deviations below the mean. The next 13.5% of users are the "*early adopters*" and are included in the area between minus two and minus one standard deviation from the mean. The next 34% are the "*early majority*" of users and are included in the area between minus one standard deviation and the population mean. Once half of the marketplace/population has adopted, the next group to adopt is the "*late majority*," making up the next 34% of users between the mean and one standard deviation above the mean. The final 16% of adopters are classified as the "*laggards*," occupying the area starting one standard deviation from the mean and continuing toward near-total adoption.

While the normal bell-shaped curve shows the adoption of an innovation over time, if the cumulative number of adopters is plotted over time, it results in an S-shaped curve superimposed over the normal distribution curve as shown in Figure 2. The S-shaped curve and the bell-shaped curve display the same data in two different ways. The S-shaped curve shows a cumulative view of adoption and the normal bell-shaped curve shows the frequency of adoption. When the cumulative number of adopters reaches a certain point, known as critical mass, continued adoption becomes self-sustaining. Critical mass represents a tipping point at which the rate of adoption rapidly increases, and the S-shaped diffusion curve takes off. The tipping point of critical mass is unique to each technology, but typically occurs at 10% to 20% adoption. Drawn on a normal distribution curve, this point is one standard deviation below the mean, which is also the transition point between the *early adopters* and the *late majority* (Rogers, 2003).

Figure 2. S-shaped diffusion curve representing the cumulative number of adopters over time overlaid onto the bell-shaped distribution frequency curve (Rogers, 2003).



The S-shaped cumulative adopter curve rises slowly at first with the *innovators*, and *early adopters*, followed by a rapid rise (rapid adoption) through the *early majority* and *late majority* categories of adopters, after which the rate of increase gradually slows as a smaller pool of adopters remain (Rogers, 2003). The intersection of the S-shaped cumulative function curve with the adopter distribution normal curve occurs at the mean (50% adoption), which is the transition point between the *early and late majority* of adopters. This is also the inflection point of the growth curve, translating it into an S-shaped cumulative diffusion curve.

METHODOLOGY

Determining current level of IPv6 adoption

The data source used to determine the current level of IPv6 adoption was Google's IPv6 Statistics ("Google IPv6 Statistics," n.d.). Google collects Internet IPv6 adoption statistics on an ongoing basis by measuring the availability of IPv6 connectivity among Google users through a measurement JavaScript that is added

to a random sample of visits to various Google web properties. The JavaScript measurement uses HTTP to fetch a URL from an IPv4-only hostname and a URL from a dual-stack hostname in random order. While previous studies have used various metrics to measure and report the level of IPv6 adoption, there is precedent in recent empirical studies for using IPv6 user statistics as a measure of IPv6 adoption (Colitti et al., 2010; Czyz et al., 2014; Nikkhah & Guerin 2016; Pickard et al., 2017). Further, a correlation analysis conducted by Pickard et al. (2018) showed Google IPv6 Statistics as a suitable proxy for global IPv6 adoption.

We analyzed Google IPv6 user data for a 114-month period beginning January 2009 and ending June 2018. An initial analysis of the Google data revealed the number of IPv6 users was consistently higher on the weekends compared to weekdays. This is attributed to more users having IPv6 access to the Internet at home rather than at work (Colitti et al., 2010; Perset 2010). The next step in the data analysis was to develop a prediction model using IPv6 user data from Google. For consistency of the month-to-month data, the Google IPv6 user data from the first Saturday of each month was used for analysis in the prediction model.

IPv6 adoption prediction model

To forecast the IPv6 adoption milestones, monthly data points from Google's IPv6 user stats were fed into SAS JMP 12 Pro and fit a growth curve to estimate the rate of adoption. An initial linear regression analysis revealed that a straight-line fit was not feasible. Subsequently, second-degree polynomial curves were fitted to the data as shown in Table 2. Data analysis began with January 2009 user stats and a quadratic formula developed using the 114 available data points. Formulas were then developed starting in annual increments to project curves for rates of change. The R^2 values indicate an excellent fit to the polynomial curves developed. Seven projection formulas were developed for each of the years' data points starting January 2009 and ending with the January 2015 starting points. Because data beginning in 2016 and later did not demonstrate statistical significance using ANOVA at $\alpha = 0.05$, it was not included in the projection calculations.

Table 2. IPv6 utilization quadratic formula models by starting date.

Data Start	Months*	Projected Google Utilization			R ²
Jan 2009	114	-29.70148	+	0.4472243*Months 0.003012*(Months - 102.5) ²	+ 0.980 27
Jan 2010	102	-30.30503	+	0.5130501*Months 0.003411*(Months - 96.5) ²	+ 0.989 59
Jan 2011	90	-29.91045	+	0.5817704*Months 0.003772*(Months - 90.5) ²	+ 0.984 60
Jan 2012	78	-28.11144	+	0.6483965*Months 0.004039*(Months - 84.5) ²	+ 0.983 03
Jan 2013	66	-24.71995	+	0.7086952*Months 0.004182*(Months - 78.5) ²	+ 0.978 95
Jan 2014	54	-21.48884	+	0.7918221*Months 0.004578*(Months - 72.5) ²	+ 0.970 74
Jan 2015	42	-15.70112	+	0.8514030*Months 0.004639*(Months - 66.5) ²	+ 0.950 99

* Count of months as of 01-July-2018

Assessing IPv6 user experience

A study on IPv6 adoption would be incomplete if it focused solely on the quantity of adoption. A qualitative analysis is also essential to assess the user experience over the IPv6 infrastructure. Poor IPv6 enablement and weak global IPv6 support can result in a poor user experience, therefore it is necessary to measure the quality of IPv6 enablement and not just the quantity of adoption. Assessing user experience involves more than measuring successful connection attempts and round-trip time (RTT). It requires measuring deeper operational metrics that reflect the user experience with services delivered over IPv6. These metrics include all components necessary for a client to access and retrieve web content on behalf of a user. This includes DNS response time, TCP connect time, and full webpage/application load time over HTTP (Popoviciu, 2016).

To assess the quality of the IPv6 user experience on the Internet, the average HTTP load times over IPv6 and over IPv4 of target websites were collected from three geographic locations within North America at 15-minute intervals over a period of 30 days. As previously mentioned, the use of HTTP load time provides an operational metric that considers Domain Name System (DNS) response time, TCP connection time, and full webpage/app load time. Eighteen user experience monitoring agents were deployed via a cloud service provider, Digital Ocean, with

six agents each in San Francisco, Toronto, and New York City. The web sites of US government agency (USGA) domains were selected as the targets for this experiment. USDA domains are mandated to operationally use native IPv6 (Kundra, 2010) providing an accessible database of a manageable number of dual-stacked domains needed for this study.

A complete list of 1,315 USGA domains was downloaded from the General Services Administration (GSA) as a .csv file (GSA, 2017). The first task was to filter out any domains no longer active and any domains not advertising IPv6 AAAA records in the global DNS. This was accomplished using a custom script to send a DNS query through Google's DNS resolver at 8.8.8.8 for each of the 1,315 domains, checking for both IPv4 A and IPv6 AAAA records. Of the 1,315 domains queried, only 600 returned IPv6 AAAA records. These 600 domains were then subjected to a reachability test from each of the monitoring agents to confirm that the sites themselves were active. Nine of the domains, even though returning AAAA records, were unreachable over either IPv4 or IPv6. The conclusion made concerning these nine domains is that the websites were no longer active, but name records had not been removed from DNS. Of the remaining 591 domains, a further 127 were unreachable by the monitoring agents over IPv6. A second reachability test to these 127 domains was conducted manually over both IPv4 and IPv6 from a web browser. This step verified that the 127 sites were in fact accessible over IPv4, but not over IPv6. In other words, the domains were active and advertised in DNS with IPv6 AAAA records but not actually reachable over IPv6. This left 464 domains verified to be reachable over IPv6.

The 464 IPv6 reachable domains were polled at 15-minute intervals for a period of thirty days from eighteen network monitoring agents deployed in three geographically distributed locations within North America: New York City, San Francisco, and Toronto. The use of multiple agent location vantage points helps avoid biases associated with an individual location (Nikkhah et al., 2011; Dhamdhere et al., 2012). 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 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 continents; the robustness in their implementation tools, including an API; and the ability to configure resource monitoring for each Droplet from the Digital Ocean dashboard.

Due to memory limitations on each Digital Ocean VM, the six agents deployed at each location monitored a subset of the 464 domains so that data from all domains was collected from each location. At each 15-minute polling interval, each agent 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 3 describes each recorded variable.

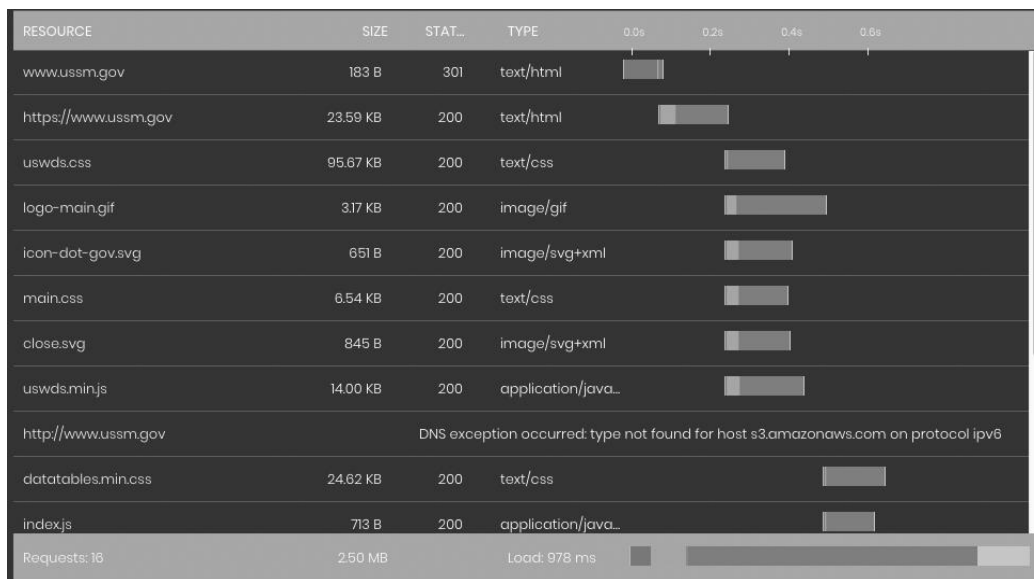
Table 3. Description of variables recorded by each network monitor

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 (McDonald, 2017).
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 (Sexton, 2015).
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 (Pingdom, n.d.).
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 (Bixby, 2010).
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 (Christensson, 2006).
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 (Rekhter et al., 2006).

At the end of the 30-day data collection period, the raw data from the 18 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 (Galletta et al., 2004). 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.

Of the 464 domains tested, 143 had at least one component that did not perform properly or timed-out when accessed over IPv6. Figure 3 shows an example HTTP waterfall for a website with failed components. A further test was conducted on each of the 464 websites using a web browser from a PC with only IPv6 access to the Internet. This test confirmed the findings from the waterfall test that components of 143 websites were not loading.

Figure 3. Example TCP waterfall showing a resource that failed to load over IPv6.



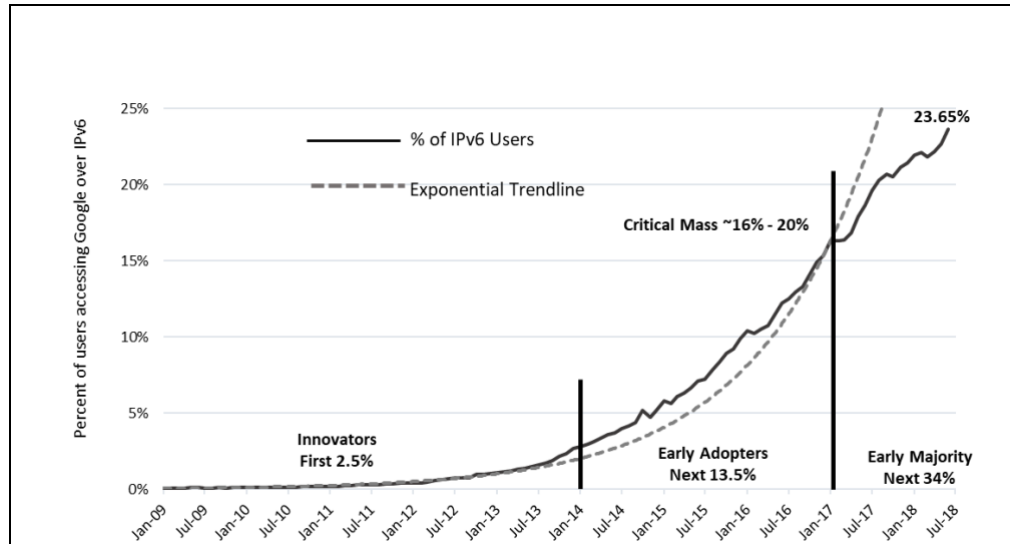
FINDINGS AND LIMITATIONS

This paper presents three significant findings. First, empirical evidence that IPv6 adoption has reached the point of critical mass on the normal distribution diffusion curve and is now solidly within Rogers' *late majority* category of adoption. Second, that there is a 95% confidence band predicting the growth rate of IPv6 adoption will continue, reaching full adoption between December 2024 and June 2026. Third, there are inconsistencies between IPv6 and IPv4 in the quality of the user experience accessing web content over the Internet. Each of these three findings are discussed next.

Critical mass of IPv6 attained

IPv6 adoption, based on Google IPv6 user data reached 23.65% on June 30 of 2018, exceeding the 16% - 20% critical mass threshold defined by Rogers (2003). The solid line in Figure 4 shows the percentage of Google IPv6 users overlaid onto the boundaries of Rogers' first three adopter categories. The data shows that it took 19 years, from the time IPv6 was standardized in 1995 until January 2014, for the first 2.5% of users, defined as *innovators*, to access Google over IPv6. Adding the next 13.5% of users, the *early adopters*, took only three years, from January 2014 to January 2017. At 23.65%, IPv6 adoption is now solidly within the *late majority* category which began at 16% of users.

Figure 4. IPv6 user adoption from January 2009 through June 2018, reaching 23.65 percent.



Future IPv6 adoption shows increase rate


Using the predictive model presented previously in Table 2, we forecast the next adopter categories of *late majority* which starts at 51% of users, *laggards* which starts at 85% of users, and finally, *full adoption* of 100% of users. The first column in Table 4 shows the data starting point fed into the prediction formula. The data endpoint was June 2018 for all seven calculations. The models yielded projection results indicating that the *late majority* of IPv6 users will begin sometime between May 2021 and February 2022, the *laggards* between December 2023 and April 2025, and full adoption between December 2024 and June 2026.

Table 4. Google IPv6 user projections made from seven data starting points yielding seven forecast start dates for each adopter category.

Data Start	Start <i>Late Majority</i> (51%)	Start <i>Laggards</i> (85%)	Full Adoption (100%)
2009	2022 Feb	2025 Apr	2026 Jun
2010	2021 Nov	2024 Nov	2025 Dec
2011	2021 Sept	2024 July	2025 Aug
2012	2021 July	2024 Apr	2025 May
2013	2021 July	2024 Mar	2025 Mar
2014	2021 May	2023 Dec	2024 Dec
2015	2021 May	2023 Dec	2024 Dec

Figure 5 shows the projected boundaries of IPv6 adoption predictions, of the author's original work, Pickard et al. (2018), compared to the predictions using current data up through June of 2018. The solid lines in the figure, noted as "original," represent the findings presented by Pickard et al. (2018) in which data collection ended March 2017. The dashed lines represent the findings that include an additional 15 months of data collected through June 2018 in this study. Of note is the narrowing of the projection boundaries that results from the updated data. The area between the dashed projection lines creates a 95% confidence band, estimating IPv6 will reach full adoption between December 2024 and June 2026.


Figure 5. Projections of IPv6 adoption milestones. Confidence band of original study by Pickard et al. (2018) shown with solid lines and confidence band of updated data collected through June 2018 shown with dashed lines.



IPv6 quality of user experience


Analysis of the quality of user experience over IPv6 began with 1,315 USGA domains. Slightly fewer than half, 45.6% or 600, had AAAA records in the global DNS. Of these, only 77.3% or 464, were reachable over IPv6. Of these 464 domain web sites, 143, or about 30%, timed out or had at least one component that failed to load, leaving 321 domain sites reachable and fully functional over IPv6. Figure 6 breaks down the quantity of USGA domains in each stage of the data analysis.

Figure 6. Number of USGA domains in each stage of data analysis.



The HTTP load time performance IPv6 reachable domain web sites are shown in Figure 7. The graph shows a comparison of IPv6 vs. IPv4 HTTP load times in terms of milliseconds with negative numbers indicating faster load times for IPv4.


Figure 7. The millisecond difference between IPv4 and IPv6 HTTP load times for 321 domains. Negative numbers indicate faster IPv4 load times than IPv6.



A review of TCP waterfall for the sites that had a slower load time over IPv6 compared to the load time over IPv4 showed that the increase in time was caused by either a longer TCP WAIT time, or a longer TCP RECEIVE time for the individual resources for the site. For sites that had longer TCP WAIT times, the increase was caused by site resources not being available via IPv6. Sites that had an increased TCP RECEIVE time were those that utilized a content delivery network (CDN) for hosting the site's static content. The total TCP RECEIVE time is increased in those cases by the client having to initiate a TCP connection to a new site, the CDN, and by downloading content. In these cases, we see the TCP RECEIVE time increases based on where the CDN is located in relation to the client requesting the resource, and by how much content needs to be retrieved. If the content were stored locally on the webserver, there is the potential for a decrease in the overall TCP RECEIVE time due to not having to create a new connection for every off-server resource that is needed. Not all CDNs are IPv6 enabled, so it is possible for the site to be available via IPv6, but that some of the resources must still be delivered via IPv4. On a dual-stack client content being provided by both IPv4 and IPv6, it is usually transparent to the user. However, when testing the sites from a PC with only an IPv6 address, those site resources that were not available via IPv6 failed to load.

Figure 8 shows the differences in HTTP load times in percentage terms with negative numbers indicating faster load times for IPv4. We found that of the 321 domains fully functional over IPv6, only ten (or 3%) had IPv6 HTTP load times within 10% or better to that of IPv4 HTTP load times. This finding contrasts with a previous study conducted by Pickard et al. (2017) measuring the user experience of Web sites from the Alexa top 100 in which 74% of the tested domains had average IPv6 HTTP load times within 10% or better to IPv4.

Figure 8. The percentage difference between IPv4 and IPv6 HTTP load times for 321 domains. Negative numbers indicate faster IPv4 load times than IPv6.



A remaining question of interest is whether there is a common property that explains the disparity in IPv6/IPv4 performance findings of industry web sites from the Alexa top 100 as reported by Pickard et al. (2017) and the IPv6/IPv4 performance findings of UGSA sites in this study. Unfortunately, our measurements did not capture any such trait or property to empirically explain this disparity. While more research is needed, we theorize that the disparity may be partly explained through an understanding of the unique technological and business or policy requirements that drove the organizations to adopt IPv6.

The sites on the UGSA list were IPv6 enabled to meet a mandate by the Office of the Federal Chief Information Officer. The mandate required agencies to “Upgrade public/external facing servers and services to operationally use native IPv6 by the end of FY 2012” (Kundra, 2010, p. 1). The mandate did not require nor did it recommend any form of qualitative analysis to monitor the quality of IPv6 enablement. In contrast, IPv6 enablement of industry Web sites included in the Alexa top 100 were likely driven by individual technology and business plans that recognize that poor IPv6 implementations can lead to poor user experience, which can have a negative impact on the brand, translating to a business cost (Popoviciu, 2016).

Limitations

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

First, Google IPv6 user data was collected from a single day of each month, the first Saturday. Measurements could be increased to include additional days and non-weekend days. Second, the scope of the user experience analysis was limited to only USDA domains. Analysis could be expanded to include non-government domains, such as industry and higher education for a more complete picture of the quality of IPv6 enablement. Third, 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

In this paper we have presented methodologies that allow us to explore and draw conclusions on several aspects of IPv6 diffusion. We have applied these methodologies to provide evidence that IPv6 adoption has reached critical mass, that adoption is predicted to continue accelerating with full adoption likely occurring between December 2024 and June 2026, and that the user experience accessing web content over IPv6 is not yet on par with IPv4. Based on these findings we conclude that IPv6 adoption is following and will continue to follow the projections of Rogers' models of innovation diffusion and will reach a point of full adoption.

However, our analysis of the user experience over the IPv6 enabled infrastructure makes clear that there is much work to be done by the Internet community to ensure that the IPv6 enabled infrastructure delivers connectivity and reachability performance that is on par or better than IPv4. This can be accomplished through user experience monitoring with end-to-end measurements that take into account all aspects of accessing applications and services. In many cases, enabling IPv6 on the external services isn't enough to ensure a robust user experience over IPv6. Additional factors such as, internal connectivity, external connectivity, and service provider support for IPv6 need to be included in any IPv6 enablement plan.

The implication of these conclusions is that if IPv6 will indeed become the next generation Internet protocol, then organizations are well advised to it as the production protocol it is and invest the time, effort, and resources necessary to ensure that IPv6 deployments are done right and that the quality of enablement is measured every step of the way.

Finally, 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. We also plan to extend our investigation into the root causes of the disparate performance behavior experienced access web content over IPv6 and IPv4.

REFERENCES

- Alghatrifi, I., & Khalid, H. B. (2018). Factors affecting the adoption of IPv6 from IPv4: A systematic mapping study. *International Journal of Current Engineering and Technology*. Vol. 8, No. 5 DOI: <https://doi.org/10.14741/ijcet/v.8.5.9>
- Beeharry, J., & Nowbutsing, B. (2016). Forecasting IPv4 Exhaustion and IPv6 Migration. *Proceedings of 2016 IEEE International Conference on Emerging Technologies and Innovative Business Practices for the Transformation of Societies*. Balaclava, pp. 336-340. doi: 10.1109/EmergiTech.2016.7737362
- 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/>.
- Christensson, P. (2006). *Traceroute Definition*. Retrieved 2018, Sep 5, from <https://techterms.com>
- Classe, A. (2003). Don't panic, but be ready to adopt IPv6. *ComputerWeekly.Com*, Retrieved from <http://www.computerweekly.com/news/2240050520/Dont-panic-but-be-ready-to-adopt-IPv6>.
- Coleman, L. (2014). Next generation internet policy in Japan, China and India. *Asia & the Pacific Policy Studies*, 1(3), 497-512.

- 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
- Czyz, J., Allman, M., Zhang, J., Iekel-Johnson, S., Osterweil, E., & Bailey, M. (2014). Measuring IPv6 Adoption. *ACM SIGCOMM Computer Communication Review*, 44(4), 87-98.
- Deering, S., & Hinden, R. (1995). Internet protocol, version 6 (IPv6) specification .
- Deering, S., & Hinden, R. (2017). Internet protocol, version 6 (IPv6) specification. RFC 8200.
- Dell, P. (2010). Two economic perspectives on the IPv6 transition. *info*, 12(4), 3-14. <https://doi.org/10.1108/14636691011057046>
- Dell, P., Kwong, C., Syamsuar, D., Francois, S., & Choy, K. (2007). A comparison of attitudes to IPv6 in three countries. In *Conference proceedings-corporate strategy and the path to a new economy*. Curtin University of Technology.
- 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 Conference on Internet Measurement Conference*, pp. 537-550.
- Domingues, M., Friacas, C., & Veiga, P. (2007). Is global IPv6 deployment on track? *Internet Research: Electronic Networking Applications and Policy*, 17(5), pp. 505-518.
- Doyle, J. (2008). European Union plans widespread IPv6 deployment by 2010. *NetworkWorld.com*, Retrieved from <http://www.networkworld.com/article/2344470/cisco-subnet/european-union-plans-widespread-ipv6-deployment-by-2010.html>
- 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.
- 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.

- Garretson, C. (2005, Aug 1.). Federal mandate to boost IPv6, but-
NetworkWorld.com, Retrieved from
<http://www.networkworld.com/article/2312998/lan-wan/federal-mandate-to-boost-ipv6--but-----.html>
- Google IPv6 statistics (n.d.). Retrieved from
<https://www.google.com/intl/en/ipv6/statistics.html>
- GSA (n.d.) General Services Administration. Gov Domains. Retrieved August 2017 from <https://catalog.data.gov/dataset/gov-domains-api-c9856>.
- Hovav, A., & Schuff, D. (2005). Global diffusion of the internet V - the changing dynamic of the internet: Early and late adopters of the IPv6 standard. *Communications of the Association for Information Systems*, 15, (1). 14
- Huston, G. (2018). Retrieved August 27, 2018, from the IPv4 Address Report: <https://ipv4.potaroo.net/>
- Internet Society (n.d.). Retrieved from
<http://www.internetsociety.org/ipv6/archive-2011-world-ipv6-day>
- Khan, R., & Sindi, H. F. (2012, June). Necessity to migrate to IPv6. In *Information Society (i-Society), 2012 International Conference* (pp. 210-215). IEEE.
- Kleinrock, L. (2010). An early history of the internet. *IEEE Communications Magazine*, 48(8), 26-36.
- Kundra, Vivek (2010, September 28). *Memorandum for chief information officers of executive departments and agencies*. Executive Office of the President, Office of Management and Budget.
- Ladid, L. (2009, June). IPv6 - The next big bail-out: Will IPv6 save the Internet? In *Proceedings of the International Conference on Computer Systems and Technologies and Workshop for PhD Students in Computing*, p.2. ACM.
- 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>.

- 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>
- Nikkhah, 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.
- 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
- NRO (2011). Retrieved August 27, 2018, from NRO <https://www.nro.net/ipv4-free-pool-depleted/>
- Perset, K. (2010). Internet addressing: Measuring deployment of IPv6. Organization for Economic Co-Operation and Development report.
- Pickard, J., Angolia, M., & Chou, T. (2018). IPv6 Diffusion on the Internet Reaches a Critical Point *Journal of Technology, Management, and Applied Engineering*, 34(1), 2-17.
- 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/>
- Pingdom (n.d.) Difference between response time and page load time. Retrieved 2 September 2018 from <https://help.pingdom.com/hc/en-us/articles>
- 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>.
- Popoviciu, C., & Dini, P. (2006, August). IPv6 as a practical solution to network management challenges. In *Computing in the Global Information Technology, 2006. ICCGI'06. International Multi-Conference* (pp. 5-5). IEEE.

- 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>>
- Rogers, E. (2003). *Diffusion of Innovations (5th ed.)*. New York, N.Y.: Simon & Schuster, Inc.
- Sexton, P. (2015, August 17). TCP Connection Time. Retrieved September 2, 2018 from <https://varvy.com/performance/tcp-connection-time.htm>
- 1
- Wang, X., & Zander, S. (2018). Extending the model of internet standards adoption: A cross country comparison of IPv6 adoption. *Information & Management*, 55(4), 450-460.
- World IPv6 Launch (n.d). Retrieved from <http://www.worldipv6launch.org/>
- Wu, J., Wang, J., & Yang, J. (2011). CNGI-CERNET2: An IPv6 deployment in China. *ACM SIGCOMM Computer Communication Review*, 41(2), 48-52.